

PSYCHOLOGY

ADVANCEANCE

DEESE JAMES DEESE JAMES DEESE JAMES

Professor of Psychology Professor of Psychology Professor of Psychology Professor of Psychology
Hopkins University Hopkins University Hopkins University Hopkins University

Second Edition Second Edition Second Edition Second Edition

Wiley Wiley Wiley Wiley

STUDENT EDUCATIONAL STUDENT EDUCATIONAL STUDENT EDUCATIONAL STUDENT EDUCATIONAL

ST 01

RG1

64452

30.3.67

154.4
D36 P



ALLAMA IQBAL LIBRARY



64452

The Psychology of Learning. Copyright © 1958 by the McGraw-Hill Book Company, Inc. Copyright, 1952, by the McGraw-Hill Book Company, Inc. All rights reserved. This book, or parts thereof, may not be reproduced in any form without permission of the publishers.

Library of Congress Catalog Card Number 58-9853

INTERNATIONAL STUDENT EDITION. *Exclusive rights by Kōgakusha Co., Ltd. for manufacture and export from Japan. This book cannot be re-exported from the country to which it is consigned by Kōgakusha Co., Ltd. or by McGraw-Hill Book Company, Inc. or any of its subsidiaries.*

PREFACE

Learning occupies one of the central positions in psychology, and an understanding of the psychology of learning should be useful to almost every serious student of psychology. The aim of the first edition of this book was to bring together for the student the divergent interests in the psychology of learning. This, likewise, is the aim of the new edition. In other respects, however, the two editions are very different. The first edition was a survey; this edition is less so. The present version treats fewer separate topics and treats these in greater detail. It is more theoretically directed than was the first edition and more selective. The more intensive coverage is achieved by eliminating topics that are of little or no fundamental importance and by eliminating peripheral topics which are generally more adequately covered in other courses. As before, however, the author has strenuously tried to resist the inclination to ride special hobbyhorses; the material is selected, but it is selected in such a way as to give the student as representative a picture as possible of the contemporary concern with the psychology of learning.

As in the first edition, the early chapters attempt to cover the basic problems of conditioning and learning. The material in these chapters deals with the acquisition of individual responses, and there is little attention given in these chapters to the interaction between responses. The basic controlling variables—reinforcement, drive, frequency, etc.—are discussed in these chapters. The middle chapters provide a bridge to problems in learning which center around interactions of learned responses. Chaining, or serial learning, and transfer of training are examples of the topics covered. The later chapters deal with retention, problem solving, and the other processes under the control of implicit and mediating activity.

The approach of the book is largely empirical, but it has been impossible (and undesirable) to avoid a theoretical bias, however slight. It is easy to

detect, and instructors will not find it difficult to present somewhat different theoretical positions without interference from the book. To the informed reader the influence of B. F. Skinner, E. C. Tolman, and the late C. L. Hull is obvious in the early chapters. Perhaps less obvious is the fact that the author is considerably indebted for the point of view in the later chapters to G. A. Miller, C. N. Cofer, C. E. Osgood, and others who have recently worked so intensively on the role of coding and mediating activity in the human intellectual processes.

Again, as was the case in the first edition, there has been some attempt to compromise between what experimental psychologists find fruitful (or easy) to investigate in the laboratory and what students see as more fundamental problems. Although there has been very little effort to treat "applied" problems as such, the relation of basic concepts in the psychology of learning to things outside the rat laboratory and the memory-drum cubicle is stressed throughout the book. The author has the fond hope that one day the psychology of learning will return to its original position as one of the basic sciences underlying educational psychology and similar applied fields. There is no doubt that some sections of this book have been influenced by that hope.

Although the book is principally addressed to the junior or senior student in psychology, graduate students in psychology and in educational psychology will find it useful. The book is relatively self-sufficient and can be studied with profit even by a student with only the slightest background in psychology. It is the author's hope, at the same time, that it will provide a foundation for further study, particularly the study of theoretical issues which can only be understood with some background in the basic empirical problems in the psychology of learning.

Numerous people have read all or portions of the manuscript. To these individuals the author gives his thanks. Many of the suggestions made by readers have been incorporated into the book; in other cases the author, perhaps rashly, has decided to let matters stand—partly because he wanted the book to reflect the contemporary scene in the psychology of learning as faithfully as possible. Special thanks are due to the author's wife, Ellin Deese, who carefully read many versions of the manuscript and who typed the larger portion of the final draft. Also, a special word of thanks is due to Clifford T. Morgan, whose very considerable editorial talents include the ability to teach others some of the skills of writing.

Permission has been granted by many publishers and individuals to reproduce illustrations or passages, and the author gratefully acknowledges

these permissions: *American Journal of Psychology*; the American Psychological Association for its several journals; Appleton-Century-Crofts, Inc.; Cambridge University Press; Ralph Gerbrands; Harvard University Press; *Journal of General Psychology*; J. B. Lippincott Company; Ronald Press Company; F. Loren Smith; *Teachers College Contributions to Education*; University of California Press; John Wiley & Sons, Inc. In each case the original source has been cited with the material reproduced.

James Deese

CONTENTS

<i>Preface</i>	v
1. Introduction	1
2. Reinforcement and Learning	15✓
3. Extinction and Partial Reinforcement	46
4. Discrimination and Learning	74
5. Motivation and Learning	98✓
✓6. Emotion and Learning	123✓
7. Serial Learning	150
8. The Learning of Skills and Conditions of Practice	181✓
✓9. Transfer of Training	213✓
10. Retention and Forgetting	236
11. Problem Solving and Thinking	273✓
✓12. Theory and Applications of Learning	301✓
<i>Name Index</i>	357
<i>Subject Index</i>	363

Conclusion

Bibliography

CHAPTER 1

INTRODUCTION

Most human behavior is learned. This fact in itself makes clear why the study of learning is one of the most important parts of psychology. It is impossible to understand the behavior of human beings and that of most animals without knowing something about the basic principles of learning. In order to have a thorough understanding of almost any problem in psychology, we should know what effect learning and different learning experiences have upon the particular problem we are studying. For example, if we were interested, as personality theorists are, in the effects of traumatic experiences during infancy or childhood upon adult personality, we should discover that we could understand this problem only to the extent that we understood the basic principles of learning. If we knew as much about learning as we should like, we could answer questions about the extent to which traumatic experiences can affect adult life and the degree to which different segments of adult behavior may be affected. In short, it is not difficult to prove the importance of learning in psychology.

It is probably more difficult to convince the student that *learning theory* is of enough general interest and importance to warrant study. Theory, even to the fairly sophisticated student of psychology, often means fussy disputes about matters of no consequence. We must admit that there is a certain justice in this view. This should not blind us, however, to the importance of theory, both in the understanding of learning itself and in its applications to the study of personality and abnormal behavior.

Theories serve several useful functions. They supply hypotheses about unsolved problems and serve to integrate and lend coherence to a scientific discipline. They help us to organize factual information and to understand

its various applications to other areas. It is usually more reasonable to formulate a hypothesis about personality development, for example, from a well-established theory of learning than from a simple factual rule. This is true because a theory frequently suggests why a particular application should be made, while the extension of a factual rule, more often than not, is simply a guess.

There is no better way to appreciate the value of theory than to read a survey of the experimental literature about some aspect of learning. If we do this, we shall find hundreds of experiments which overlap, duplicate, and frequently contradict one another. With no theory to guide us, we should find it difficult to tell important from unimportant experiments. If we were to memorize the contents of such a survey, we should have at our command a mass of highly specific information limited to narrow examples of behavior. Most general principles are established through theory of some sort. Therefore a theory can be used to lift us beyond the confines of isolated experiments into an awareness of general principles about learning and behavior.

Unfortunately there are very few theoretical principles in the psychology of learning that are both unequivocal and widely accepted, because very few existing theories are supported by a wide variety of evidence. We must therefore be cautious in the extent to which we apply theory. It is necessary to strike a balance between the simple factual information at our disposal and the general principles developed or invented to explain and extend it. Since complete devotion to a given theory is likely to be dangerous given our present psychological ignorance, we should be cautious in accepting a universal set of explanatory principles that can be applied to all problems in learning. Rather, we shall study some limited aspects of certain theories and their applications to specific problems. Some of the grand sweep of a really systematic theory is lost in this way, but, considering the limitations of our factual knowledge, this approach is more realistic.

SOME METHODOLOGICAL BACKGROUND

Before we can discuss the factual evidence and theoretical principles of the psychology of learning, we need to establish some ground rules about finding evidence and discovering or inventing theoretical principles. We shall now look briefly at the nature of experiments and the logic of scientific method as they apply to some of the problems we shall examine later.

Experiment and Observation

A scientific observation occurs when one looks at an event and records it in such a way that any other competent observer, given the same conditions, would make the same report. An experiment, in the simplest sense, is an observation made under conditions that are, to some extent, of the observer's own devising. In an experiment the investigator can actually change and control the events that occur. Thus, watching squirrels retrieve buried nuts from the ground could be a scientific observation, but to turn this into an experiment, we should have to control some of the conditions of observation (such as the place where the nuts are buried).

Most of the factual information in the psychology of learning has been gathered in experiments, because, whenever possible, investigators prefer to experiment rather than simply to observe. This preference is well founded; it is always easier to infer *causal* relations when the events under study are controlled. The kind and extent of the control in the observation limit the adequacy and generality of the causal inferences the investigator may make. If only one condition is systematically varied, we can be a good deal more certain of the cause than if all conditions were allowed to vary haphazardly. Since we are usually interested in making causal inferences from our observations, it will be worthwhile to glance at some of the rules for conducting experiments.

Independent and dependent variables. There is a particular way in which events are controlled in an experiment. In the simplest kind of experiment, all conditions are held constant but one, and this is allowed to vary in known ways. For example, we may want to study the effects of different kinds of rewards upon the speed with which monkeys can learn a new habit. To do this, we should systematically vary the nature of the reward given to different monkeys and try to keep constant in all monkeys the effects of motivation, previous history, testing conditions, etc. Since we had controlled all the variables but reward, we should be safe in making inferences from our results about the effects of quality of reward upon speed of learning.

In this example, kind of reward is an *independent variable* because it is independently controlled and varied by the experimenter. The changes in behavior that result from changes in reward provide the *dependent variable*. Speed of learning, in this example, is the dependent variable because changes in it depend upon changes in the independent variable, reward.

Independent variables, then, are the conditions that are systematically varied by the experimenter; the dependent variables are the consequent changes in behavior.

More than one independent or dependent variable. In the example just mentioned, all the variables but kind of reward were controlled, either by being held constant or by being allowed to vary randomly from animal to animal. In some experiments, however, we may use more than one independent variable and measure more than one dependent variable.

Experiments with more than one independent variable are quite common. Investigators frequently use this method because they suspect that the two independent variables somehow modify each other's influence on the dependent variable. In these circumstances, we speak of the result as an *interaction* between the two variables. For example, reward and motivation interact in their influence on behavior. Under very low or under very high motivation, animals will frequently not discriminate among rewards of slightly different quality. But under moderate motivation—a mild degree of hunger, for example—animals will discriminate among qualitatively different rewards. Thus, it is easy to see that the effect of reward on behavior is *modified* by the effect of motivation on behavior.

Sometimes, too, we may use more than one dependent variable in an experiment. We may be interested in the influence of reward upon both speed of learning and general activity. In this case we obtain information that we should not get by making two separate experiments, one on speed of learning and one on general activity level. The additional information is not quite in the same form as that achieved in the case of two independent variables; usually it comes in the form of a *correlation* between the two dependent variables. Speed of learning and general activity may be so highly correlated, for example, under the influence of different kinds of reward or motivation, that we suspect there is some underlying general factor that accounts for the changes in both speed of learning and general activity.

Functional relationships. When the dependent variable shows change as a result of variations in the independent variable, it is said that the dependent variable is a *function* of the independent variable. In the example we have discussed, speed of learning is a function of quality of reward. Functional relationships are very important in the psychology of learning, since many of the results are quantitative and can sometimes be expressed as simple mathematical functions.

More frequently a functional relationship between the independent and dependent variable is presented in the form of a graph. Graphs are useful

because they show at a glance the nature of a functional relationship. In graphic presentation the independent variable is always placed on the horizontal axis (the abscissa) and the dependent variable on the vertical axis (the ordinate). A good example of such a graph is Figure 4 on page 21. Here the magnitude of a conditioned response (dependent variable) is plotted against the number of trials of training (independent variable).

Definitions in the Psychology of Learning

The term reward that we have been using is a common word in our vocabulary. It means something to everyone, but, unfortunately, nearly everyone interprets it in a slightly different way. If a term is to be useful to the experimental psychologist, it must be capable of exact definition; that is, it must refer eventually to a measurable fact—a concrete object or condition. Many words that refer to psychologically important events are also used in an everyday sense. The common definitions of these terms, even when explicitly set forth in a good dictionary, depend upon casual observations, tradition, and the opinion of authorities. Of course, there can be as many definitions of this sort as there are dictionaries.

Operational definitions. In order to make our psychological language as explicit as possible, we must redefine many familiar words as well as introduce some new words. Then we shall have to give these words clear *operational* meaning. An operational definition of a term, in the simplest case, is a description of the operations performed by an experimenter or observer and the changes in behavior that result. Desire, reward, motives, etc., can all be given surprisingly precise definitions if described in this way. A hungry rat, according to the operational definition commonly agreed upon by psychologists, is one that has been deprived of food for a period of time after having been on an accustomed feeding schedule.

By the same token, an operational definition of the term reward usually states that it is some change in the environment (presenting an animal with food, for example) that increases the likelihood of the animal repeating the behavior that immediately preceded the reward. Notice that there is no mention of “something that the animal likes” or “something that it wants to get”; rather, the definition is explicitly stated in terms of things that can always be observed and agreed upon. We may all disagree about what animals “like” (unless we operationally define “like”), but most of us can agree that sometimes animals repeat what they have done before because it was followed by a reward. As a matter of fact, this is the fundamental aspect of voluntary learned behavior.

Hypothetical constructs. Of course, not all the terms used by psychologists are given direct operational definitions. Someone, for example, may say that "learning is an increase in the excitatory potential for a given habit." Excitatory potential in this definition is a hypothetical construct; it is not something that can be directly observed. Excitatory potential is a theoretical term invented by a psychologist to account for certain changes in behavior. Psychology abounds with these hypothetical constructs; as a matter of fact, the word learning itself is usually used in the sense of a hypothetical construct. Hypothetical constructs are useful (1) when they can be related to things that are operationally defined and (2) when they have wider application than an operational definition.

The second requirement of a hypothetical construct calls for some explanation. Let us take the example of reward again. Reward is occasionally used as a hypothetical construct rather than an operationally defined term (though more usually the term reinforcement is substituted for reward). In such a case it may be defined by saying that it is a "slight decrement in the level of neural excitability." If we define reward as a hypothetical construct, we must show that such a definition is useful. It would be useful if it could serve to relate our operational definition of a reward (in terms of change of behavior) to some other conditions. In this case we might argue that the hypothetical construct serves as a guess as to *why* the behavior changes. If this is so, we ought to be able to draw a number of implications from the hypothetical definition of reward and put them to experimental test. If the construct stands up under test, we might conclude that we had an "explanation" of the operational fact of the effects of reward on behavior. Thus, we see that hypothetical constructs are potentially very useful as explanations.

Unfortunately it is sometimes very difficult to put hypothetical constructs to a test (the above example is particularly bad in this respect), and consequently psychology abounds in constructs that are *potentially* explanatory but are only really so if we are willing to accept a great many dubious assumptions. Nevertheless, we still find it necessary to make use of constructs in the following discussion, even though we know that they are not always well grounded in experimental fact.

Some psychologists make a distinction between hypothetical constructs and intervening variables (MacCorquodale and Meehl, 1948). In this distinction, hypothetical constructs are theoretical terms which *cannot* be entirely reduced to collections of operational definitions, while intervening variables are, in a sense, abstractions which can be. Intervening variables are terms that serve to relate a number of operationally defined independent

and dependent variables. One example of an intervening variable that we shall see much of in this book is *strength of response*. Strength of response is not a dependent variable, but it determines what will happen to dependent variables. Thus, if we suspected that quality of reward affected strength of response, we should have no way of measuring this effect directly. We could, however, measure the effect of quality of reward upon the likelihood of future occurrences of a response and upon the magnitude, or size, of future responses. We should expect both likelihood and magnitude to increase with preferred rewards. We can then use the single term strength of response to cover both changes in probability of occurrence and size of responses.

If we make a hypothetical construct out of the notion of strength of response, we add something to its definition that cannot be directly referred to a collection of observations. We might say, for example, that strength of response is ultimately defined in terms of the configuration of protein molecules on the nerve membrane. It is unlikely that we shall ever directly observe such configurations of molecules, but such a definition could have important implications for our interpretation of experimental facts.

This distinction between hypothetical constructs and intervening variables may seem abstract and, indeed, it is. When dealing with the intricacies of psychological theory, however, it is sometimes helpful, and we shall have occasion to use it in many connections. While intervening variables are useful in limited psychological theory, whenever a theoretical psychologist tries to relate the theoretical terms of psychology to neurophysiological or physical events and theories, he makes use of hypothetical constructs.

If we keep in mind the distinction between hypothetical constructs and intervening variables, as well as the distinction between these and operational definitions, we shall find it much easier to keep fact and theory separate and to evaluate theory against available factual evidence.

The remainder of this chapter discusses an important operational distinction between two types of behavior and two types of learning. Since there is a question as to whether or not this distinction involves theoretical as well as operational matters, the importance of keeping operational and theoretical terms separate should be apparent at the outset.

TWO TYPES OF LEARNING

Learning to typewrite is different from learning to ski, if for no other reason than that different movements and muscles are involved. Learning to

typewrite and learning to ski, however, are both examples of the same processes of learning at work. In both cases, initial attempts are corrected in response to information the learner receives about their success. Most psychologists, however, are now convinced that not all learning involves the same basic processes, though they do not agree about how many processes there are. Some psychologists have suggested several types; Tolman (1949), for example, suggested six different learning processes. While few theorists agree with this proposition, it is clear that one must make an operational distinction between at least two types of learning when talking about the behavior of complicated organisms like dogs and men. Perhaps, if one is concerned with more primitive organisms, it is necessary to add even more types (Humphrey, 1933). At least two types of learning are necessary, however, and a great many investigators at different times have distinguished between them.¹ Since most of these distinctions have elements in common, it is apparent that many psychologists, when distinguishing two types of learning, have been driving at some basic principles. Let us examine two cases of learning and see if we can abstract some of these principles.

Two Examples of Learning

Conditioned salivation in the dog. Every student of psychology is familiar with the famous experiments on conditioning made by Ivan P. Pavlov, the Russian physiologist. Because Pavlov's experiments are both well known and fundamental to the problem we are discussing, we can use one of them to illustrate one of our two types of learning.

In this experiment (Pavlov, 1927), small openings were made in the cheeks of dogs, so that the duct of the parotid salivary gland could be directed outward; this allowed the saliva to be collected and measured as it was secreted. The dog was placed in a soundproofed room, which had a small window that permitted the experimenters to watch the dog. A tuning fork was sounded in the room and, a few seconds later, small amounts of powdered meat were presented for the dog to eat. The tuning fork, of course, did not produce saliva, but the sight of meat and its presence in the mouth did. After a few pairings of the tuning fork and the meat powder, the tuning fork was presented to the dog *without* the meat powder. It was apparent that the tuning fork had acquired the ability to elicit the salivary response which was initially limited to the meat powder.

¹ For example, see Skinner (1938), Schlosberg (1937), Hilgard and Marquis (1940), Mowrer (1947), Konorski (1950), Stephens (1942).

Pavlov referred to the meat powder as an *unconditioned stimulus*. An *unconditioned stimulus* is any stimulus that has the ability to elicit a response without prior training.² Thus, when meat powder is placed in the mouth of a dog, a salivary reflex is immediately produced. This is called an *unconditioned response*. The tuning fork was referred to by Pavlov as a *conditioned stimulus*. Conditioned stimuli, in Pavlov's experiments, are those which initially do not elicit the response under study but which come to do so by being paired with the unconditioned stimulus. Likewise, the salivary response produced by the tuning fork is called a *conditioned response* because it is a learned response conditional upon the previous pairing of stimuli.

This basic experiment of Pavlov's has had an enormous influence upon the psychology of learning, and the terms unconditioned and conditioned stimuli and unconditioned and conditioned responses are part of the basic vocabulary of the psychology of learning. Let us now look, however, at an example of learning that is quite different from the one studied by Pavlov.

Conditioned bar pressing in the rat. Our second example is perhaps not quite so familiar, but it is equally important. This time we shall deal with the behavior of the laboratory rat, and as in the case of Pavlov's experiment, this example will serve as a prototype for an important class of learning processes.

A white rat is placed in a 12-inch-square box that is relatively sound-proof. At one end of this box is a small lever, or bar, that projects from the wall. This bar is connected to an automatic recording device and to a magazine filled with small pellets of rat food. When the bar is pressed, a pellet of food is dropped into the box from the magazine. The essentials of this device, usually called a Skinner box, can be seen in Figure 1 (Skinner, 1938).

If the rat is hungry when it is placed in the box, it will explore the box quite readily. It will sniff the air, paw the walls of the box, bite here and there. Eventually, because there is little else to do, the rat will manipulate the lever. Sooner or later it will depress the bar enough to release a pellet of food. The rat may not discover the food immediately, but when it does, it will certainly eat the food. It is also fairly certain that the rat will press

² The ability of the unconditioned stimulus to elicit a response may be because of learning that took place before the animal came to the laboratory. Operationally, however, the "unconditioned" ability of a stimulus refers to the state of things at the beginning of an experiment. Obviously, if one wishes to state that a particular unconditioned stimulus has an innate eliciting tendency, he must carefully control the life history of the animal.

the lever a second time. This time, perhaps, it will discover the pellet of food immediately. At this point the behavior of the rat will change dramatically. Instead of resuming its casual exploration of the box, it will now proceed to press the lever repeatedly. As a matter of fact, if it is hungry enough, it will press the lever at a rate determined only by the time it takes to put its nose down and eat the food.

It is apparent that the rat has learned or has been conditioned to press the bar. Is this example of learning like the one studied in Pavlov's laboratory? To answer this question, let us see if we can make use of Pavlov's

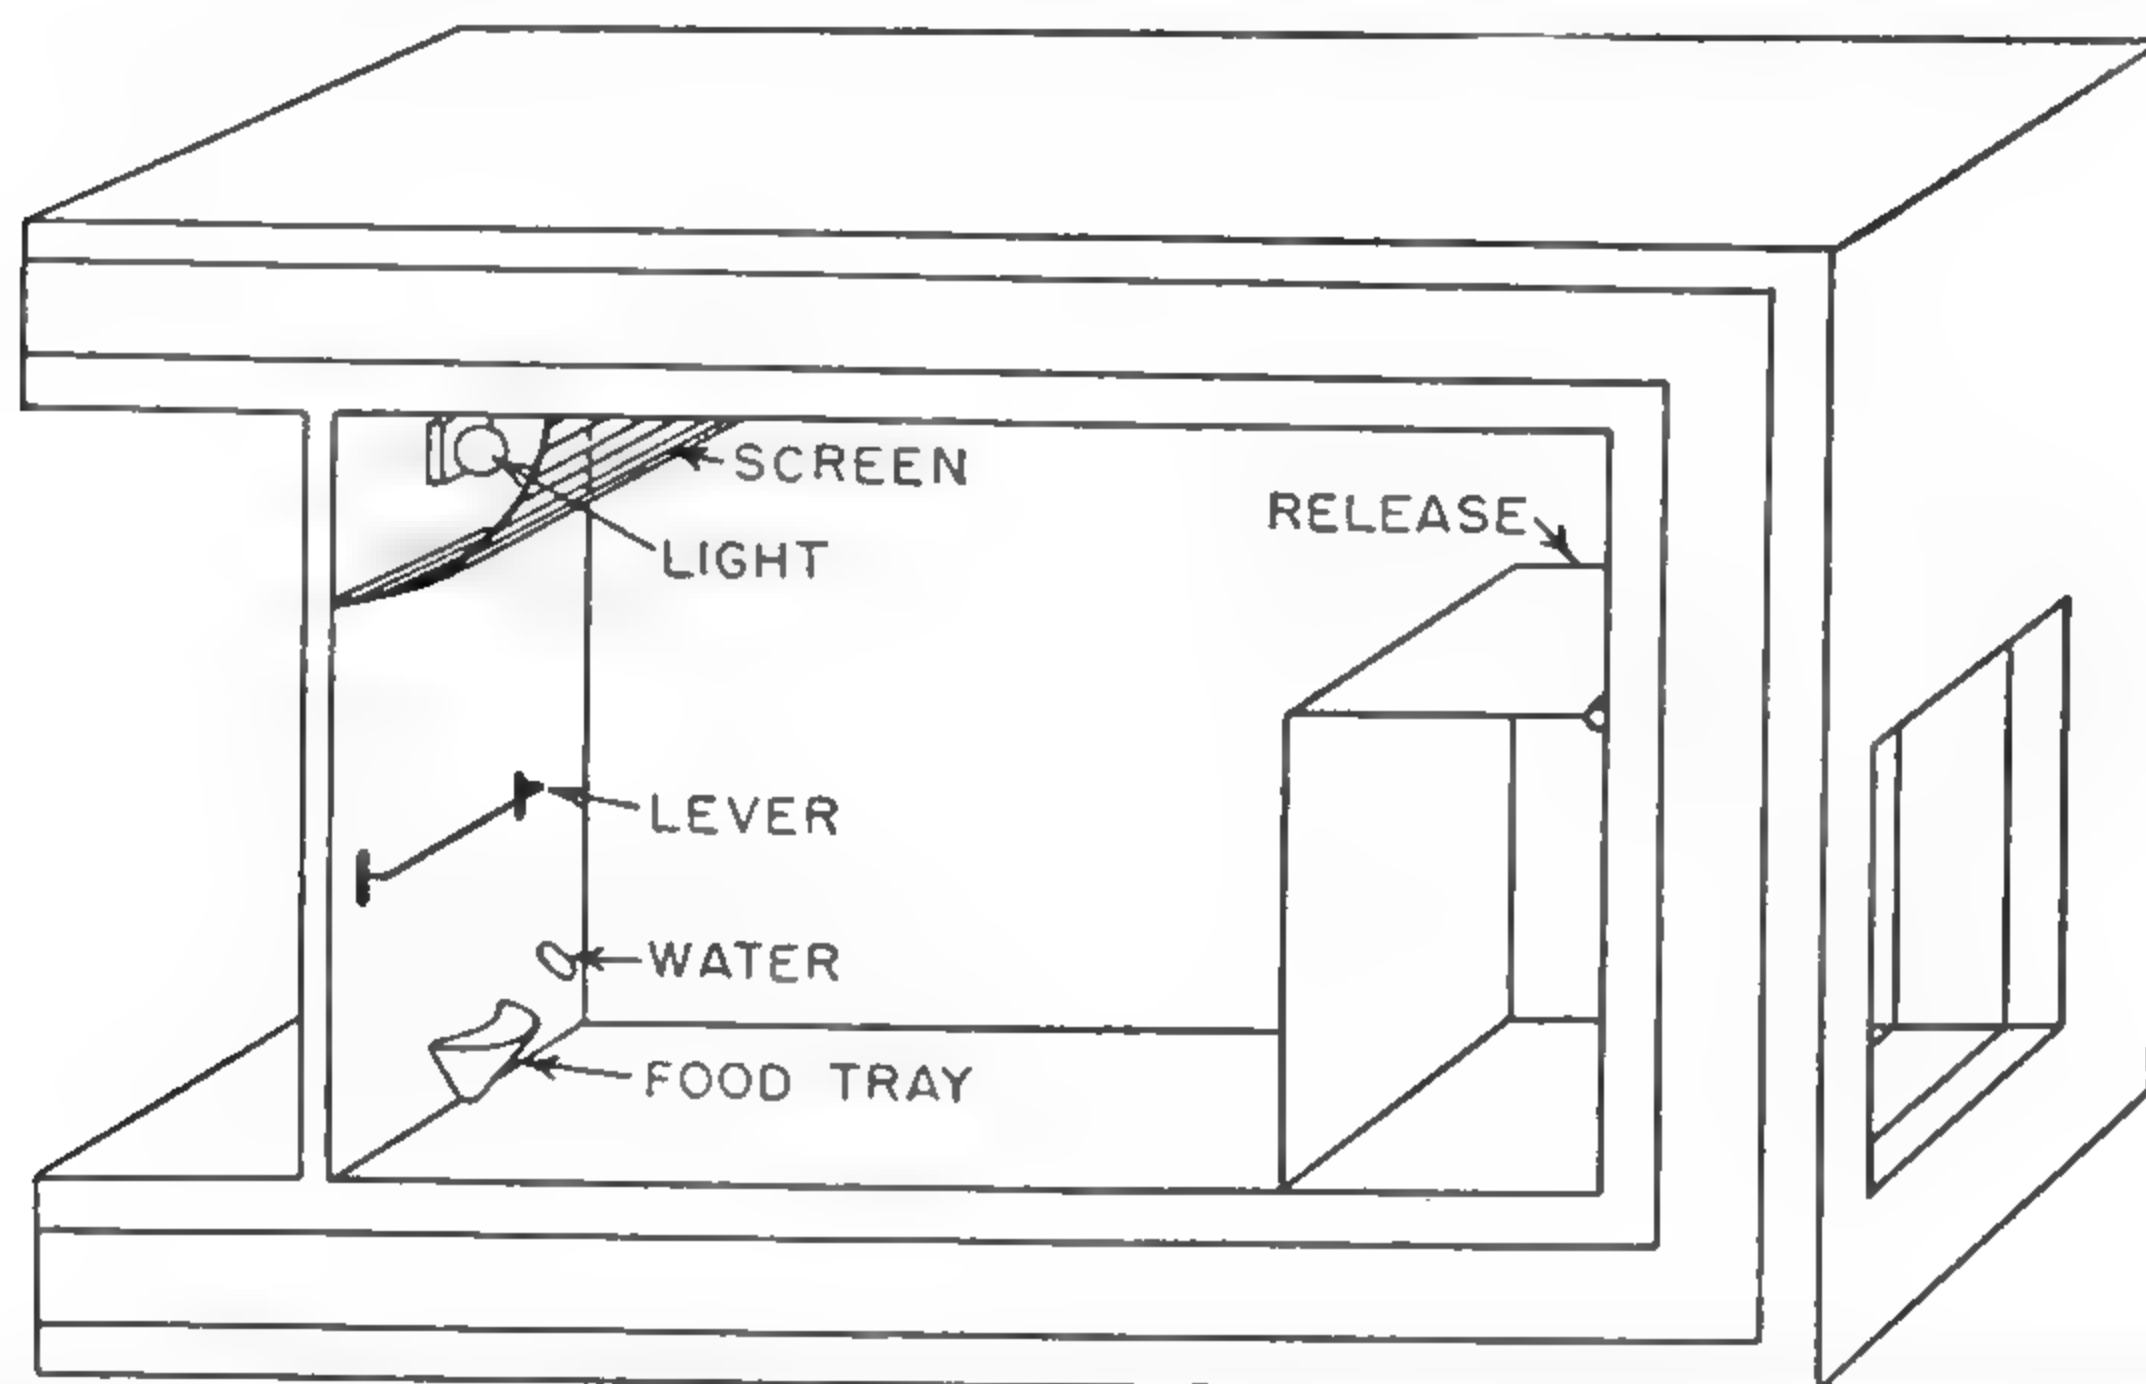
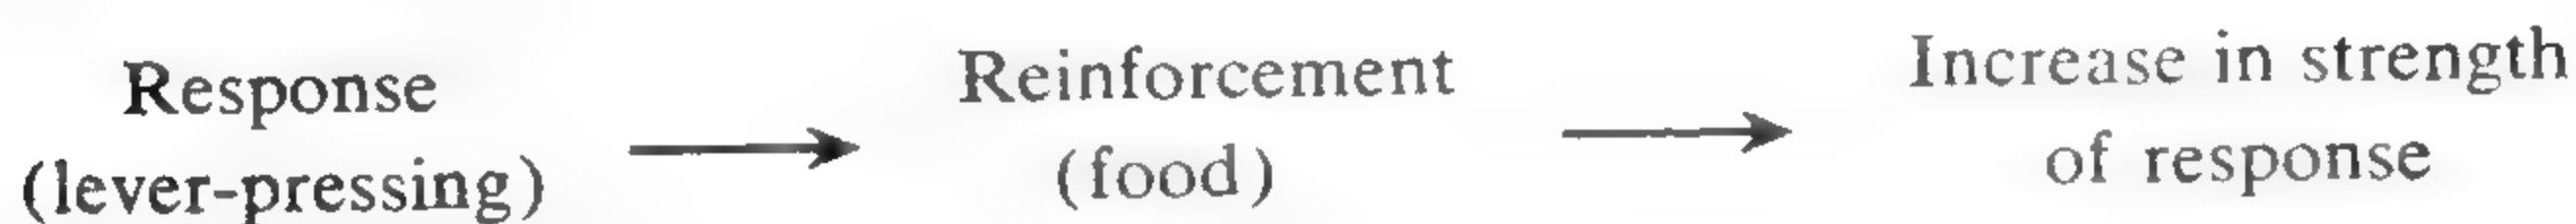


Figure 1. An example of the Skinner box. One side has been cut away to show the interior. When the animal depresses the lever, a device automatically delivers a pellet of food into the box. (Skinner, 1938.)

terms in describing this example. Our first problem arises when we try to identify the unconditioned stimulus. Now, it is clear that the unconditioned response must be pressing the bar during the course of exploration. What was the stimulus that elicited this response? We can guess, but we can never know with any certainty what the stimulus actually was. It might have been the sight of the bar or the tactual stimulation obtained by feeling along the wall. What, then, is the stimulus to which the bar pressing is conditioned (the conditioned stimulus)? In other words, what leads the rat to press the bar after it has "discovered the connection" between the lever and food? Again, we are reduced to a guess.

All we do know is that we have a response (pressing the bar) that dramatically changes its character after the rat has been rewarded a few times

for performing the response. In short, we have a response which increases in likelihood as the result of reward. The more common psychological term for reward in this situation is reinforcement. We may diagram the sequence of events this way:



The reinforcement produces the increase in the strength of the bar-pressing response.

A comparison of the two examples. There are two important differences between the two experiments we have just discussed. In the salivary-conditioning example, the response is directly under the control of the investigator. Provided that the animal is in normal condition, the salivary response can be elicited simply by providing the appropriate stimulus. In the case of the lever-pressing example, however, the investigator cannot directly elicit lever pressing; rather, he must wait until it occurs spontaneously.

The second difference is in the action of the reinforcement. Pavlov used the terms reinforcement and unconditioned stimulus interchangeably. This usage has caused some confusion, though in a sense it is justified. The unconditioned stimulus is a reinforcer in Pavlov's experiment because without it no conditioning would occur; it *reinforces* conditioning. The two terms, reinforcement and unconditioned stimulus, are not synonymous, however, for we have just learned that we cannot use the term unconditioned stimulus to apply to the reinforcement in the bar-pressing example. The reinforcement in this case *follows* the response, and we do not know what the stimulus is that *leads* to the response. Thus another difference between Pavlov's experiment and the bar-pressing experiment is that in Pavlov's experiment the unconditioned stimulus is also the reinforcement, whereas in the lever-pressing experiment no unconditioned stimulus can be identified, though a reinforcement which follows the response can usually be easily identified. In Pavlov's experiment the animal is reinforced whether or not it gives a conditioned response, while in the bar-pressing example the animal is reinforced only when a response occurs. This is a very important difference.

Principles Distinguishing the Two Types of Learning

We have looked at examples of two types of learning. Let us now turn to the basic principles that distinguish these two types; we have already

been introduced to some of these principles by comparing the two examples with each other, but in this section we shall make the principles more explicit and general.

Operant versus respondent behavior. In a long series of papers and in his book *The Behavior of Organisms* (1938), Skinner has pointed out that we have not only two basic types of learning but also two basic kinds of responses. These two types of responses Skinner called respondents and operants. Respondents are acts or behavior directly elicited by stimuli; they obey the classical physiological laws of reflexes, and, indeed, the terms respondent and reflex are identical for all practical purposes. Operants, on the other hand, are acts or behavior for which there are no readily observable external stimuli. Examples of operant behavior are the familiar ones of spontaneous activity in animals and men—monkeys pulling at bananas and infants crawling about the floor. Although we cannot identify a prior stimulus producing such behavior, we do not assume that the behavior is truly spontaneous in the sense of being without antecedents, but we recognize the fact that the stimuli producing these acts are beyond our immediate control. They may be under the control of some autonomously active center in the nervous system or under the control of some internal stimuli, such as those coming from the sense organs within the body, but in any event we usually cannot observe, much less produce, these stimuli.

Skinner has pointed out that respondents are most characteristic of the autonomic nervous system, while operants are characteristic of the skeletal nervous system. The correlation is not perfect, for the familiar knee jerk, which is skeletal, is a respondent. The correlation of these two types of behavior with different parts of the peripheral nervous system is an important one, however, because of the role of the autonomic nervous system in emotion. As one might expect, respondents are especially important in emotional learning.

Instrumental versus classical conditioning. We still have not examined the principles distinguishing two different ways of learning. The operant-respondent distinction generally refers to types of behavior, not types of learning. Hilgard and Marquis (1940) have most clearly distinguished between two types of learning, which they call classical and instrumental conditioning. The essential element defining classical conditioning is that the reinforcement be independent of the occurrence of a conditioned response. It was because Pavlov's experiments were of this type that Hilgard and Marquis characterized it as classical conditioning. The essential condition defining instrumental conditioning is that the reinforcement be de-

pendent upon the occurrence of the conditioned response. This is the case in the bar-pressing example, and indeed in any case in which an animal is "trained to do a trick." Most of our attempts to teach children and animals are of this type; we make a reward or reinforcement contingent upon the behavior we want the subject to learn.

Two types of behavior and two types of learning. In the example from Pavlov, the response was a respondent and the conditioning classical. In the example of the rat pressing the lever, the response was an operant and the conditioning instrumental. Do these correlations between respondent-classical conditioning and instrumental-operant conditioning always hold? Probably not, though in practice classical conditioning nearly always occurs with a respondent, and instrumental conditioning with an operant. There are, however, apparently mixed examples, and while we shall not deal with these at present, they will become important when we come to theoretical matters.

There is considerable dispute as to whether the operational distinctions between two types of behavior and between two types of learning are basic theoretical distinctions as well. Are these differences limited to our techniques for investigating behavior, or do they reflect some fundamental differences in the processes governing behavior and learning? There is a distinction of hypothetical constructs that corresponds to the operational distinctions we have been discussing, and cases which are borderline in operational distinctions become very important in deciding whether or not theoretical distinction should be made. These are problems, however, that we can defer for a while.


CONDITIONS CONTROLLING SIMPLE LEARNING

We have seen two examples of elementary kinds of learning. Such examples are usually called conditioning. There is no special way in which conditioning can be clearly differentiated from learning, but the term usually refers to cases of learning in which simple responses are acquired under simple or restricted conditions. Some investigators (Mowrer, 1947) would like to limit the term conditioning to the Pavlovian or classical variety, but such a restriction is difficult to maintain, since, as we shall see, some more complicated examples of learning involve elements of both classical and instrumental conditioning. Furthermore, the two types obey many of the same laws. In this book, the word conditioning will be used to refer to isolated responses learned under simple conditions which preclude previous or

concurrent learning that would either help or interfere with the learning of the response under study. These elementary examples are extremely interesting to the psychologist, because in them he can see at work basic processes that may be obscured in more complicated cases.

Variables that determine the nature of conditioning. In the next few chapters, we shall look at some of the variables that determine the course of simple learning. These generally operate in more complicated examples of learning as well, but usually in ways that make them difficult to study. These conditions include such things as the kind and level of motivation, the relevance of the response and the reinforcement to that motivation, the scheduling of reinforcement (whether or not reinforcements are given for every response), time relationships between stimuli and between stimuli and responses, etc.

Theories about conditioning. As we discuss the way in which these various conditions determine the course of conditioning, we shall also examine the role theories play in interpreting or predicting the results of experiments using these variables. Theories make predictions, based on the properties of hypothetical constructs, about the effects of such things as motivation upon learning and behavior. In so far as it is possible, we shall try to evaluate the success of various theories in making such predictions; in practice this is difficult to do, since theories are frequently not stated clearly enough to make unequivocal predictions about behavior. In one sense good theories are easy to prove wrong, since they make clear and unequivocal predictions that are relatively easy to test; poor theories sometimes endure because they are so vaguely stated that it is impossible to prove them wrong.



CHAPTER 2

REINFORCEMENT AND LEARNING

Reinforcement is one of the basic concepts in the psychology of learning. In practice it often means the same thing as the term reward. Reward is too general in meaning for our purposes, however; also, it is debatable whether its ordinary meaning applies to reinforcement in the case of classical conditioning. Therefore we shall not make much use of the term reward; instead we shall substitute the more exact term reinforcement.

AN OPERATIONAL DEFINITION OF REINFORCEMENT

Recall, for a moment, the example of instrumental conditioning given in the first chapter. In this example, a laboratory rat learned to press a bar, and this response was reinforced with a pellet of food. The food was clearly responsible for an increase in the rate at which the rat pressed the bar, for we know that without it the rat would only occasionally have pressed the bar. Having earned a few pellets of food in this way, however, it pressed the bar just about as rapidly as it could, pausing between responses only to eat the pellets.

If, instead of following the bar-pressing response with a pellet of food, we had used the sound of a buzzer, would the rat have learned to press the bar? Probably not, our common experience with rewards would tell us (though, it should be remarked, sometimes our common experience can lead us to the wrong conclusions). Here, however, our common sense would be right; in this situation food is a reinforcement and a buzzer is not. Food is reinforcing because it can produce an increase in the frequency of the response that it follows. Obviously, because the buzzer cannot produce such an increase, it is not a reinforcement.

A definition of reinforcement. With this example in mind, we can now look at a formal definition of reinforcement. *A reinforcement is any stimulus event that will increase or maintain the strength of a response or stimulus-response connection associated with it.* A reinforcement in the case of operant behavior is a stimulus that follows the occurrence of the response and rewards it; a reinforcement in respondent behavior is the stimulus that elicits the unconditioned reflex.

The one feature of this definition that has not yet been made explicit is the phrase *strength of response*. Strength of response is properly an intervening variable (see Chapter 1), something that cannot be directly observed. Rather we infer the strength of a response from some aspect of the response that we can measure. Thus, in order to specify the formal definition of reinforcement completely, we need to say exactly what dependent variables can be used to indicate the strength of response.

Strength of response. The most important indicator of strength of response is a measure ordinarily called probability of response. This is the likelihood that a response will occur in a given unit of time or when a particular stimulus is presented. Thus if a response occurs frequently in a given unit of time, it is a strong one. For rats in the Skinner box, a strong tendency to press the lever would be indicated by a rate of 10 responses per minute, while a weak tendency would be 1 response per minute.

Some investigators prefer to limit the measurement of strength of response to probability (Skinner, 1938, 1950); and still other investigators try to reduce all other possible indicators of strength of response to some form of probability of occurrence (Estes, 1950; Mueller, 1950). In practice, however, many other indicators are used. Two of these are reaction time or latency (Graham and Gagne, 1940) and amplitude or magnitude of response (Hovland, 1937). For the time being we shall not be too concerned with the problems created by recognizing several indicators of strength of response.

Testing the effects of reinforcement. In the case of operant behavior and instrumental conditioning, the methods of assessing the effects of reinforcement are quite simple. Usually the response before reinforcement is very low in strength; the probability of occurrence is very small. After a few reinforcements, however, the probability of occurrence per unit time changes rapidly. Figure 2 shows this effect taking place. In Figure 2 the cumulative frequency of responses is plotted against time. Thus when the frequency of responses per unit time is low, as it is at the outset of the experiment, the responses cumulate very slowly; the result is a nearly hori-

zontal line with occasional bumps which show the occurrence of responses. When the response becomes conditioned, however, the response occurs much more frequently, so that the line goes up rapidly. The actual learning itself, in this example, is shown by the change in the slope of the curve.

In classical conditioning the situation is usually more complicated. As explained in the first chapter, the reinforcing stimulus in the classical conditioning of a respondent is the unconditioned stimulus. Thus, the stimulus that elicits the response in the first place also provides the reinforcement. The conditioned or learned response, however, occurs to the conditioned stimulus, and because the conditioned stimulus usually precedes the unconditioned stimulus, the conditioned response usually "anticipates" the unconditioned stimulus. This is not necessarily the case, however, for it is

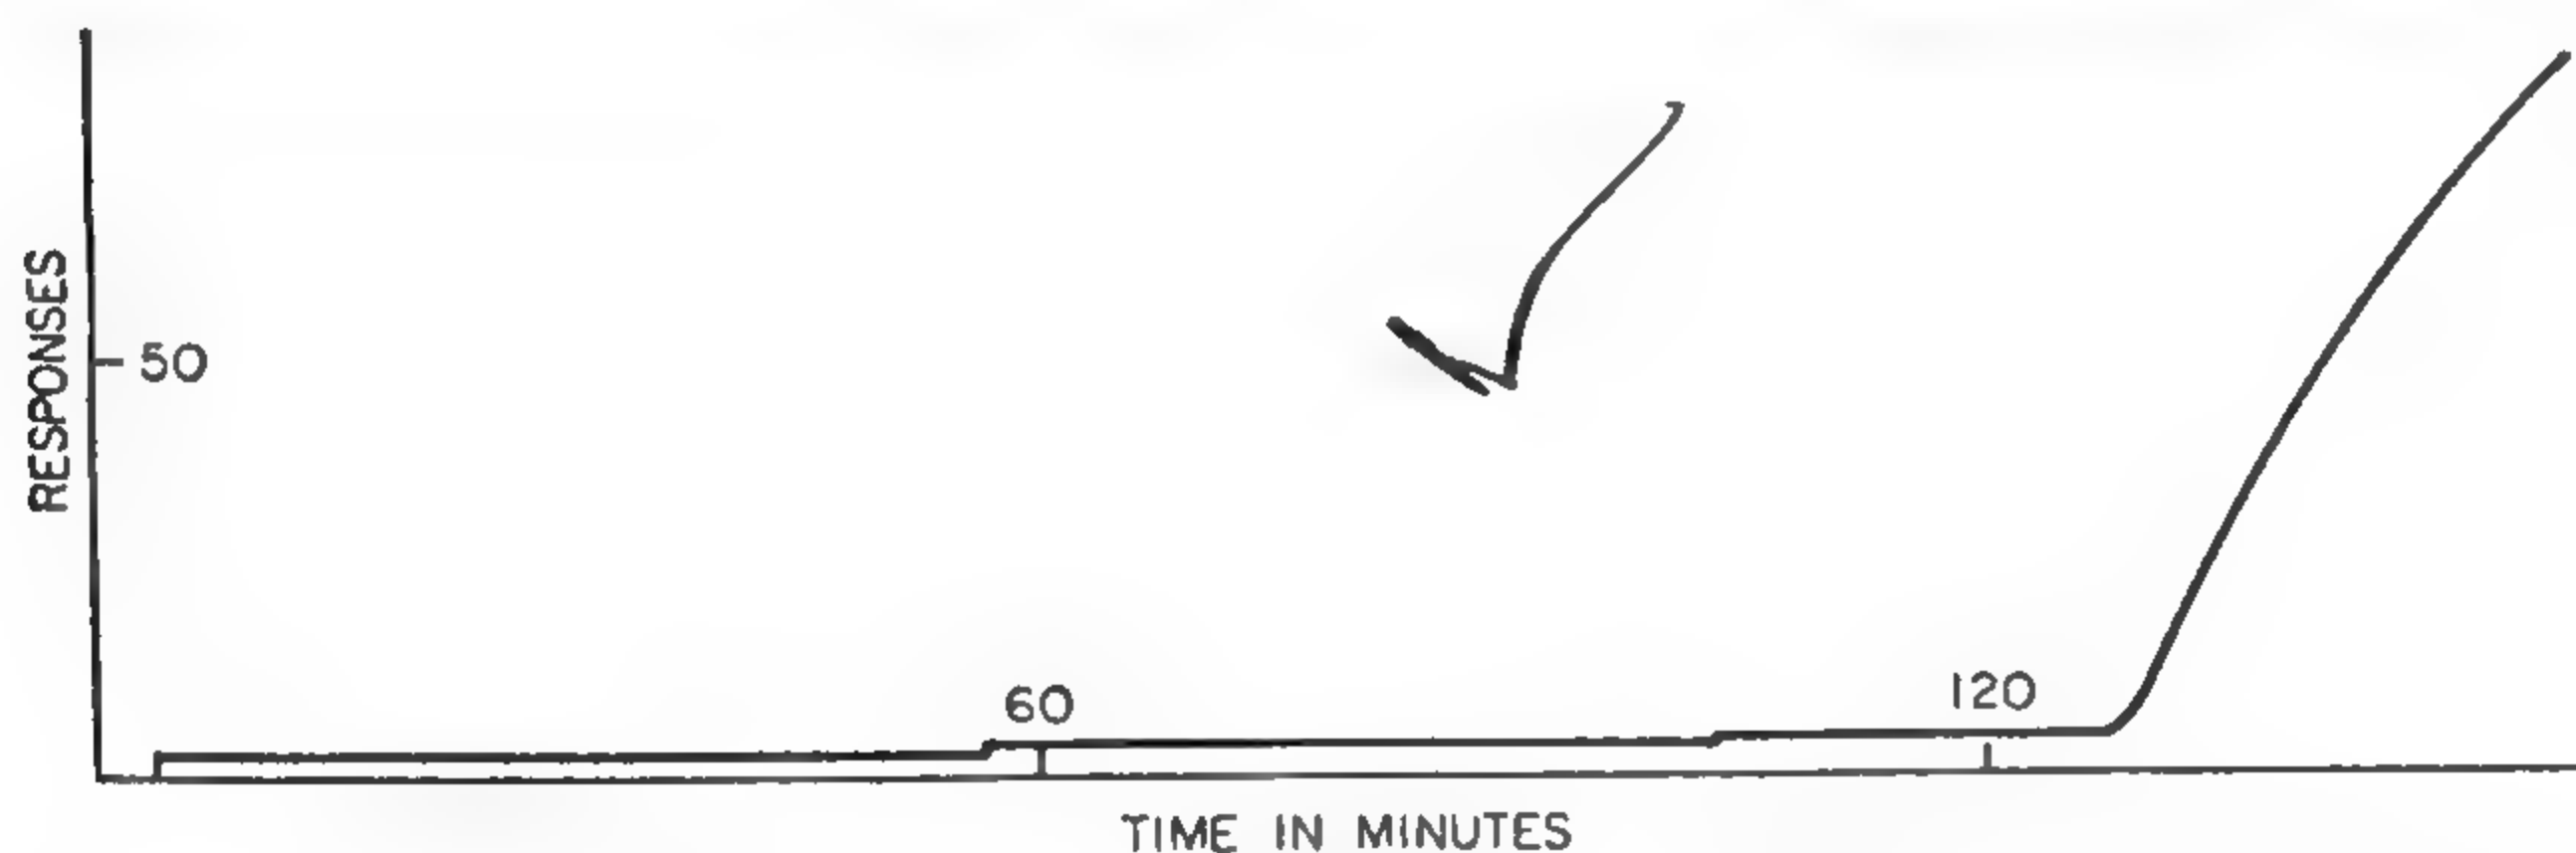


Figure 2. A cumulative curve of instrumental conditioning. All responses were reinforced. There is no evidence of learning following the first three responses, but the fourth is followed by a rapid change in rate of responding. (Skinner, 1938.)

possible to condition a respondent when the unconditioned and conditioned stimuli are simultaneous. If they occur simultaneously, then we must use special "test" trials to find out whether conditioning has taken place. Without these trials we should have no way of knowing whether a particular response was the result of the conditioned or of the unconditioned stimulus. If we make the two stimuli simultaneous and then test for a conditioned response by omitting the unconditioned stimulus on some trial, we can determine whether or not we have a conditioned response. According to the definition of reinforcement, however, prior association of the conditioned stimulus-response connection with the unconditioned stimulus or reinforcement is necessary, so that if we omit the unconditioned stimulus on too many successive test trials the conditioned response will disappear. Incidentally, in the example of salivary conditioning, the indicator of strength of response is usually one of magnitude of response—the amount of saliva given on any trial.

Trans-situational reinforcers. The definition of reinforcement given above is open to the charge of being circular—it is only a definition, not a law. There are laws about reinforcement, however, and one of them is that most reinforcers for operant behavior are trans-situational (Meehl, 1950). This means that most reinforcements will strengthen *all or most* learnable responses in a given species. In other words, if a particular stimulus event is reinforcing for a bar-pressing response in the rat, it ought to be reinforcing for other responses, such as running a simple maze. Turning out a light is reinforcing for the laboratory white rat. This reinforcement can be used to teach a rat to press a lever that turns out the light or to run a simple maze that has a dark box at the end.

Respondents are trans-situational in the sense that one can use almost any conditioned stimulus for a given response, but because of the elicited character of the respondent, the unconditioned stimulus cannot be used as a reinforcer for more than a few closely related respondents. Thus respondents are not trans-situational so far as responses are concerned.

Meehl (1950) proposes that all reinforcers (at least for operants) are trans-situational and that every increase in strength of response or learning involves the use of a trans-situational reinforcer. This is probably not exactly true, for later in this book there are examples of operant reinforcers that are specific to certain responses, but it is probably true that most of the reinforcers of everyday life are trans-situational. This means that one can use almost any reinforcement to teach almost any response; the success of the reinforcement for a particular response will not depend upon that particular response-reinforcement connection but upon the general adequacy of the reinforcement. If some stimulus is not reinforcing for bar pressing, it will not be reinforcing for learning a maze.

CONDITIONS OF REINFORCEMENT

// Let us now look at some of the experiments that tell us about the conditions under which reinforcement occurs and how these conditions modify the effects of reinforcement. These experiments are important because they tell us how behavior is determined by reinforcement when other variables either are held constant or interact with reinforcement. Since reinforcement is such an essential mechanism in the behavior of organisms, one can look upon this section as a fundamental introduction to the psychology of learning.

and the other of which does not. These other examples would show characteristics much like those of the two cases we have cited. Response strength would receive its greatest increment with the first or one of the first reinforcements, and with successive reinforcements the increase in response strength would become progressively smaller. This uniformity in simple learning curves has led psychologists to believe that there is a fundamental learning process that can be mathematically expressed as a theoretical learning curve. Let us glance at the properties psychologists have proposed for such a curve and the equation that governs it.

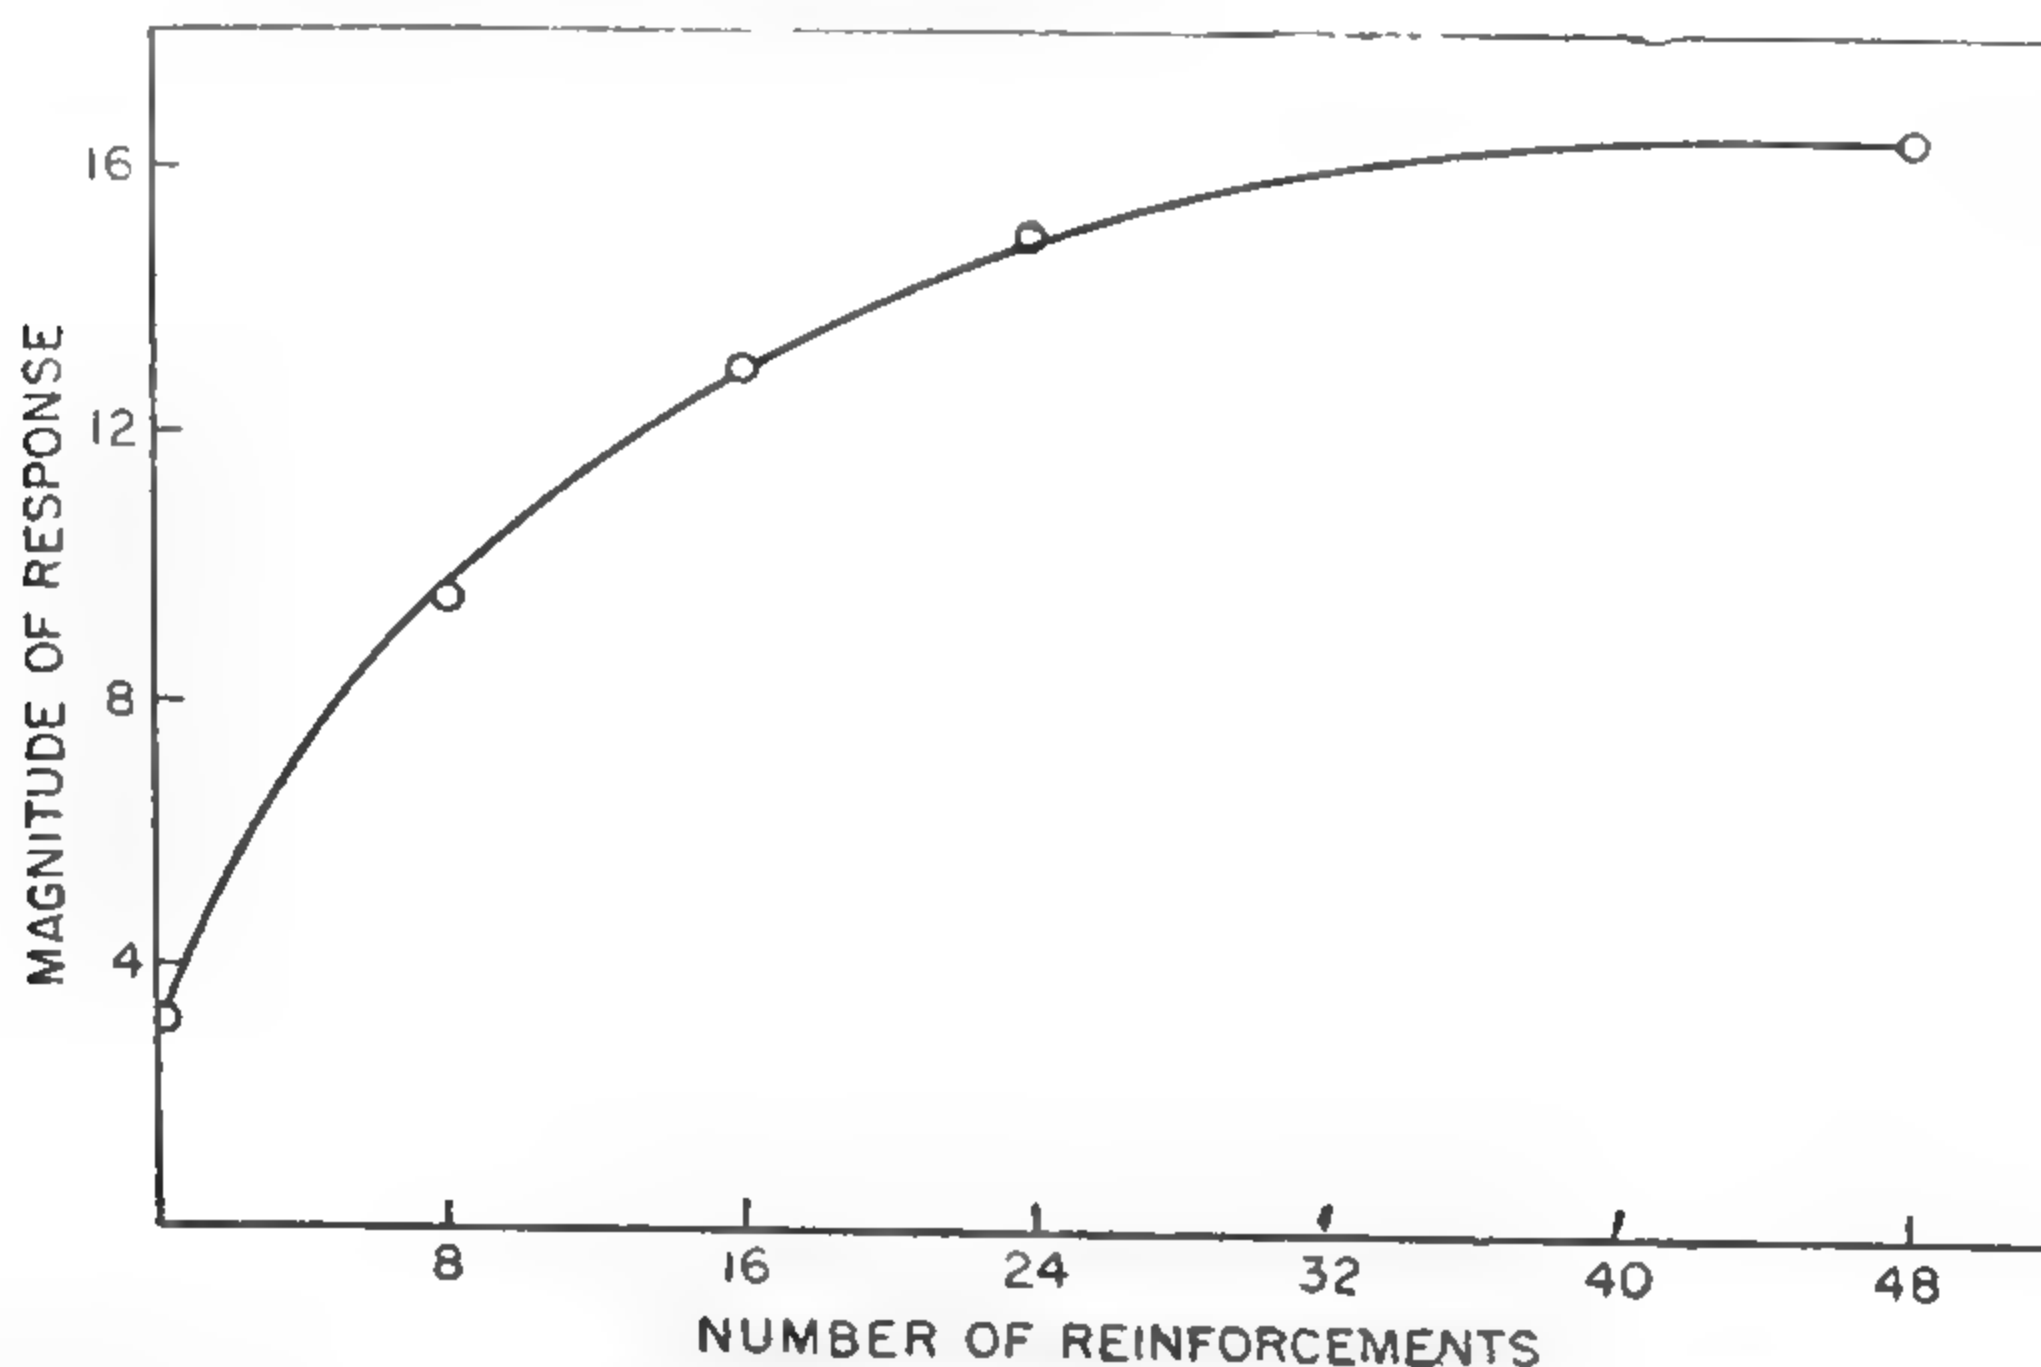


Figure 4. Amplitude of the galvanic skin response as a function of the number of reinforcements. (Hovland, 1937a.)

A theoretical curve. The exact shape of the fundamental learning curve—the one relating response strength to number of reinforcements—has occupied the interest of many theoretical psychologists. Investigators have tried to predict the shape of simple learning curves from mathematical assumptions and then to apply these predictions, in turn, to more complicated relationships in psychology. We shall now examine one of the most widely accepted of these attempts to discover the basic properties of the fundamental learning curve.

Here are three of these basic properties that should be expressed in any mathematical or theoretical learning equation:

1. Strength of response is an increasing function of the number of reinforcements or reinforced trials.

2. Strength of response increases only to some limit; beyond this there is no further increase unless the conditions of the experiment change. This limit or *asymptote* may vary with different situations and different species of animals, but with the same animals and the same situation, it is a constant.

3. Each successive reinforcement adds smaller and smaller amounts of response strength.

With these characteristics in mind let us consider how strength of response R grows with number of reinforcements N . With each increase in number of reinforcements, there is some increase in strength of response. We can put this in the form of a ratio,

$$\frac{\Delta R}{\Delta N}$$

where Δ means increase or increment. This expresses characteristic 1. Characteristic 2 says that strength of response can increase only to some limit. Let us call this limit M . Characteristic 3 suggests that as strength of response gets close to the value of M , the increases in strength of response grow smaller. Thus as the difference $M - R$ becomes smaller, ΔR becomes smaller. All of this can be put in the form of an equation:

$$\frac{\Delta R}{\Delta N} = k(M - R)$$

This equation says that the increment in strength of response is proportional to the *difference between the maximum strength possible and the amount already accumulated*. As this difference becomes zero, increase in response strength becomes zero—no further increase is possible.¹

The above equation may be written as a differential equation and integrated. This yields

$$R = M - Me^{-kN}$$

To those with little or no mathematics, this equation may not mean much. Suffice it to say that it describes learning curves of the sort presented earlier very well.

This equation, or something very close to it, has frequently been a basic assumption in theories of learning (Hull, 1943; Estes, 1950; Estes and Burke, 1953; Bush and Mosteller, 1953). Sometimes it is elaborated to provide for the possibility that the first reinforcement may not always produce the greatest increase in strength, that there may be a warm-up period,

¹ Some investigators have proposed quite different properties for basic learning curves. See, for example, Gulliksen (1934).

so that the largest increments come after the first few reinforcements. In this case the curve will be S-shaped. Despite the possibility of minor modifications, the ideas presented above tell us about the basic elements that go into theoretical learning curves. Sometimes these elements are described in terms of hypothetical neural events or *samples* of stimuli conditioned on each trial, but in any case the characteristics outlined remain the same.

Amount and Type of Reinforcement and Strength of Response

The fundamental fact about learning is that behavior changes progressively as a function of the number of reinforcements. This relationship, however, is modified by a number of conditions of reinforcement, such as the amount and quality of the reinforcing agent. In the case of food reinforcement, for example, a large amount of food results in a greater strength of response than a small amount (Grindley, 1929). In itself this is trivial, for it is limited to the kind of animals used in the experiment (chickens, in this case), the particular response, and the kind of reinforcement. Other experiments, however, give us grounds for believing that this relationship is generally true for simple conditioned operants and food reinforcement. Guttman (1953), for example, finds that it is true of the lever-pressing response in rats for various concentrations of sugar and that it is also true of peanuts used as reinforcement for problem solving in monkeys (Harlow and Meyer, 1952).

Even if we use a large experimental literature to establish a general law about strength of response and amount of reinforcement, it seems rather trivial, if for no other reason than that it appears to be obvious. There is something about this relationship that is not quite so obvious, however, and that is the *way* in which amount of reinforcement affects behavior. The exact way in which this happens involves an important distinction in the psychology of learning.

The distinction between learning and performance. For some time psychologists have been aware that there could be a difference between what an animal learns and what it does. It is possible, for example, that a rat could learn that pressing the lever in the Skinner box leads to food but would not show that it had learned, perhaps because it was afraid of the noise the food-delivery apparatus makes. Or a rat may know the location of food in a maze, but because it is not hungry, it may not show any particular disposition to go to the food. Thus, the fact that an animal does not perform very well in a learning test does not mean that the animal has not learned what it should. It is possible, however, to devise special tests that

make it possible to tell whether an animal has learned or not, even when its performance is poor.

This distinction between learning and performance is important to the question of how amount of reinforcement affects behavior, for there is a good deal of evidence to suggest that amount of reinforcement affects performance but not necessarily learning.

Two investigators (Crespi, 1942, 1944; Zeaman, 1949) have attempted to separate the effects of amount of reinforcement on learning from the effects on performance by (1) training rats to run a simple runway, such as the one described earlier in this chapter, under various amounts of food reinforcement, and then (2) switching the amount of reinforcement when the rats had reached a stable level of performance. The results of these studies are well defined. The amount of reinforcement has no effect on the *rate* at which animals approach the final level of performance, but it does affect the final level they achieve. In other words, after learning is complete, animals that have received different amounts of reinforcement do perform at different levels, but they do not learn at different rates. Thus, the limit of response strength M is affected by the amount of reinforcement, but the rate at which this limit is reached, k , is not. Furthermore, if one trains a rat with one amount of reinforcement and then suddenly shifts to a higher amount, the change in performance is almost instantaneous—the rat does not have to learn to appreciate the new amount of reinforcement; it produces an instantaneous change in the rat's behavior.

Both of these studies used a very simple response. The only thing the rats had to do was to run down a runway to get food placed at the end. One investigator (Crespi, 1944) measured speed of running, and the other investigator (Zeaman, 1949) measured latent period. Other investigators who have studied more complicated examples of learning have reported that amount of reinforcement makes no difference at all. For example, in teaching rats to discriminate between a black goal box and a white goal box by placing food in one and not in the other, it was discovered that the amount of food makes no difference in either learning or performance (Reynolds, 1949). In another example amount of food reinforcement made little difference in performance in a complicated maze, except when no food at all was given (Furchtgott and Rubin, 1953). There is a good reason why these examples of more complicated behavior, when compared with the simple running response, show little or no effects of amount of reinforcement. In the runway case, the measures of behavior are more a reflection of performance than of learning. There is very little to learn on the runway. In

the more complicated examples, because the animals have relatively difficult problems to learn, the measures of behavior more nearly reflect learning. Because the behavior reflects more learning than performance, amount of reinforcement does not appear to change behavior very much. If no food at all is given, however, performance deteriorates badly.

✕There are several other studies in the experimental literature (Maher and Wickens, 1954; Reynolds, 1950a, 1950b) which support the conclusion that learning is usually independent of amount of reinforcement while performance may not be. There is, however, a fundamental problem that is raised by one of these studies.

In this study the investigator tried training rats in a maze with no reinforcement at all. It would seem reasonable, by extending the principle that we have already stated, to say that the absence of reinforcement affects performance, not learning. This is fundamental because it suggests that reinforcement is necessary for performance but not for learning itself. This problem is so basic and so controversial that it has not yet been fully settled. For years investigators have been trying to find ways to decide whether reinforcement affects only what an animal does or both what it does and what it learns. One of the ways in which they have tried to reach a decision in this matter is to compare learning under some reinforcement with learning under no reinforcement. Such experiments are called latent-learning experiments.

Latent learning. The latent-learning experiments grew out of the theories of E. C. Tolman (1932, etc.). One of Tolman's principal ideas was that reinforcement—or reward—affected performance but had little or nothing to do with learning. The latent-learning experiments were designed to test this idea, and one of the earliest was made by Blodgett (1929).

Blodgett allowed one group of hungry rats to explore a maze that did not have food in the goal box. This group of rats, therefore, was not reinforced, argued Blodgett (though this argument is suspect, as we shall see). A second group of rats explored a maze with food in the goal box, and therefore it received the conventional reinforcement.

There was a difference between the behavior of these two groups. The group which was reinforced with food seemed to learn on schedule, but the unreinforced rats showed little improvement. After seven days, the rats in this group were entering almost as many blind alleys as they had on the first day. At this time, however, Blodgett put food into the goal box for these rats. Their performance suddenly and dramatically improved, so that almost immediately they made as few errors as did the previously reinforced rats.

These results can be seen in Figure 5. They show us that the rats which did not receive the food as reinforcement were able to profit from their experience in the maze, though the results of such experience did not show up until a reward was provided. Because the learning that apparently took place during the unreinforced trials was not evident until food was introduced, the learning was said to be "latent."

Many investigators have repeated this or similar experiments (Thistlethwaite, 1951), but not all investigators have been able to find what Blodgett found. A careful repetition of Blodgett's study (Reynolds, 1945b),

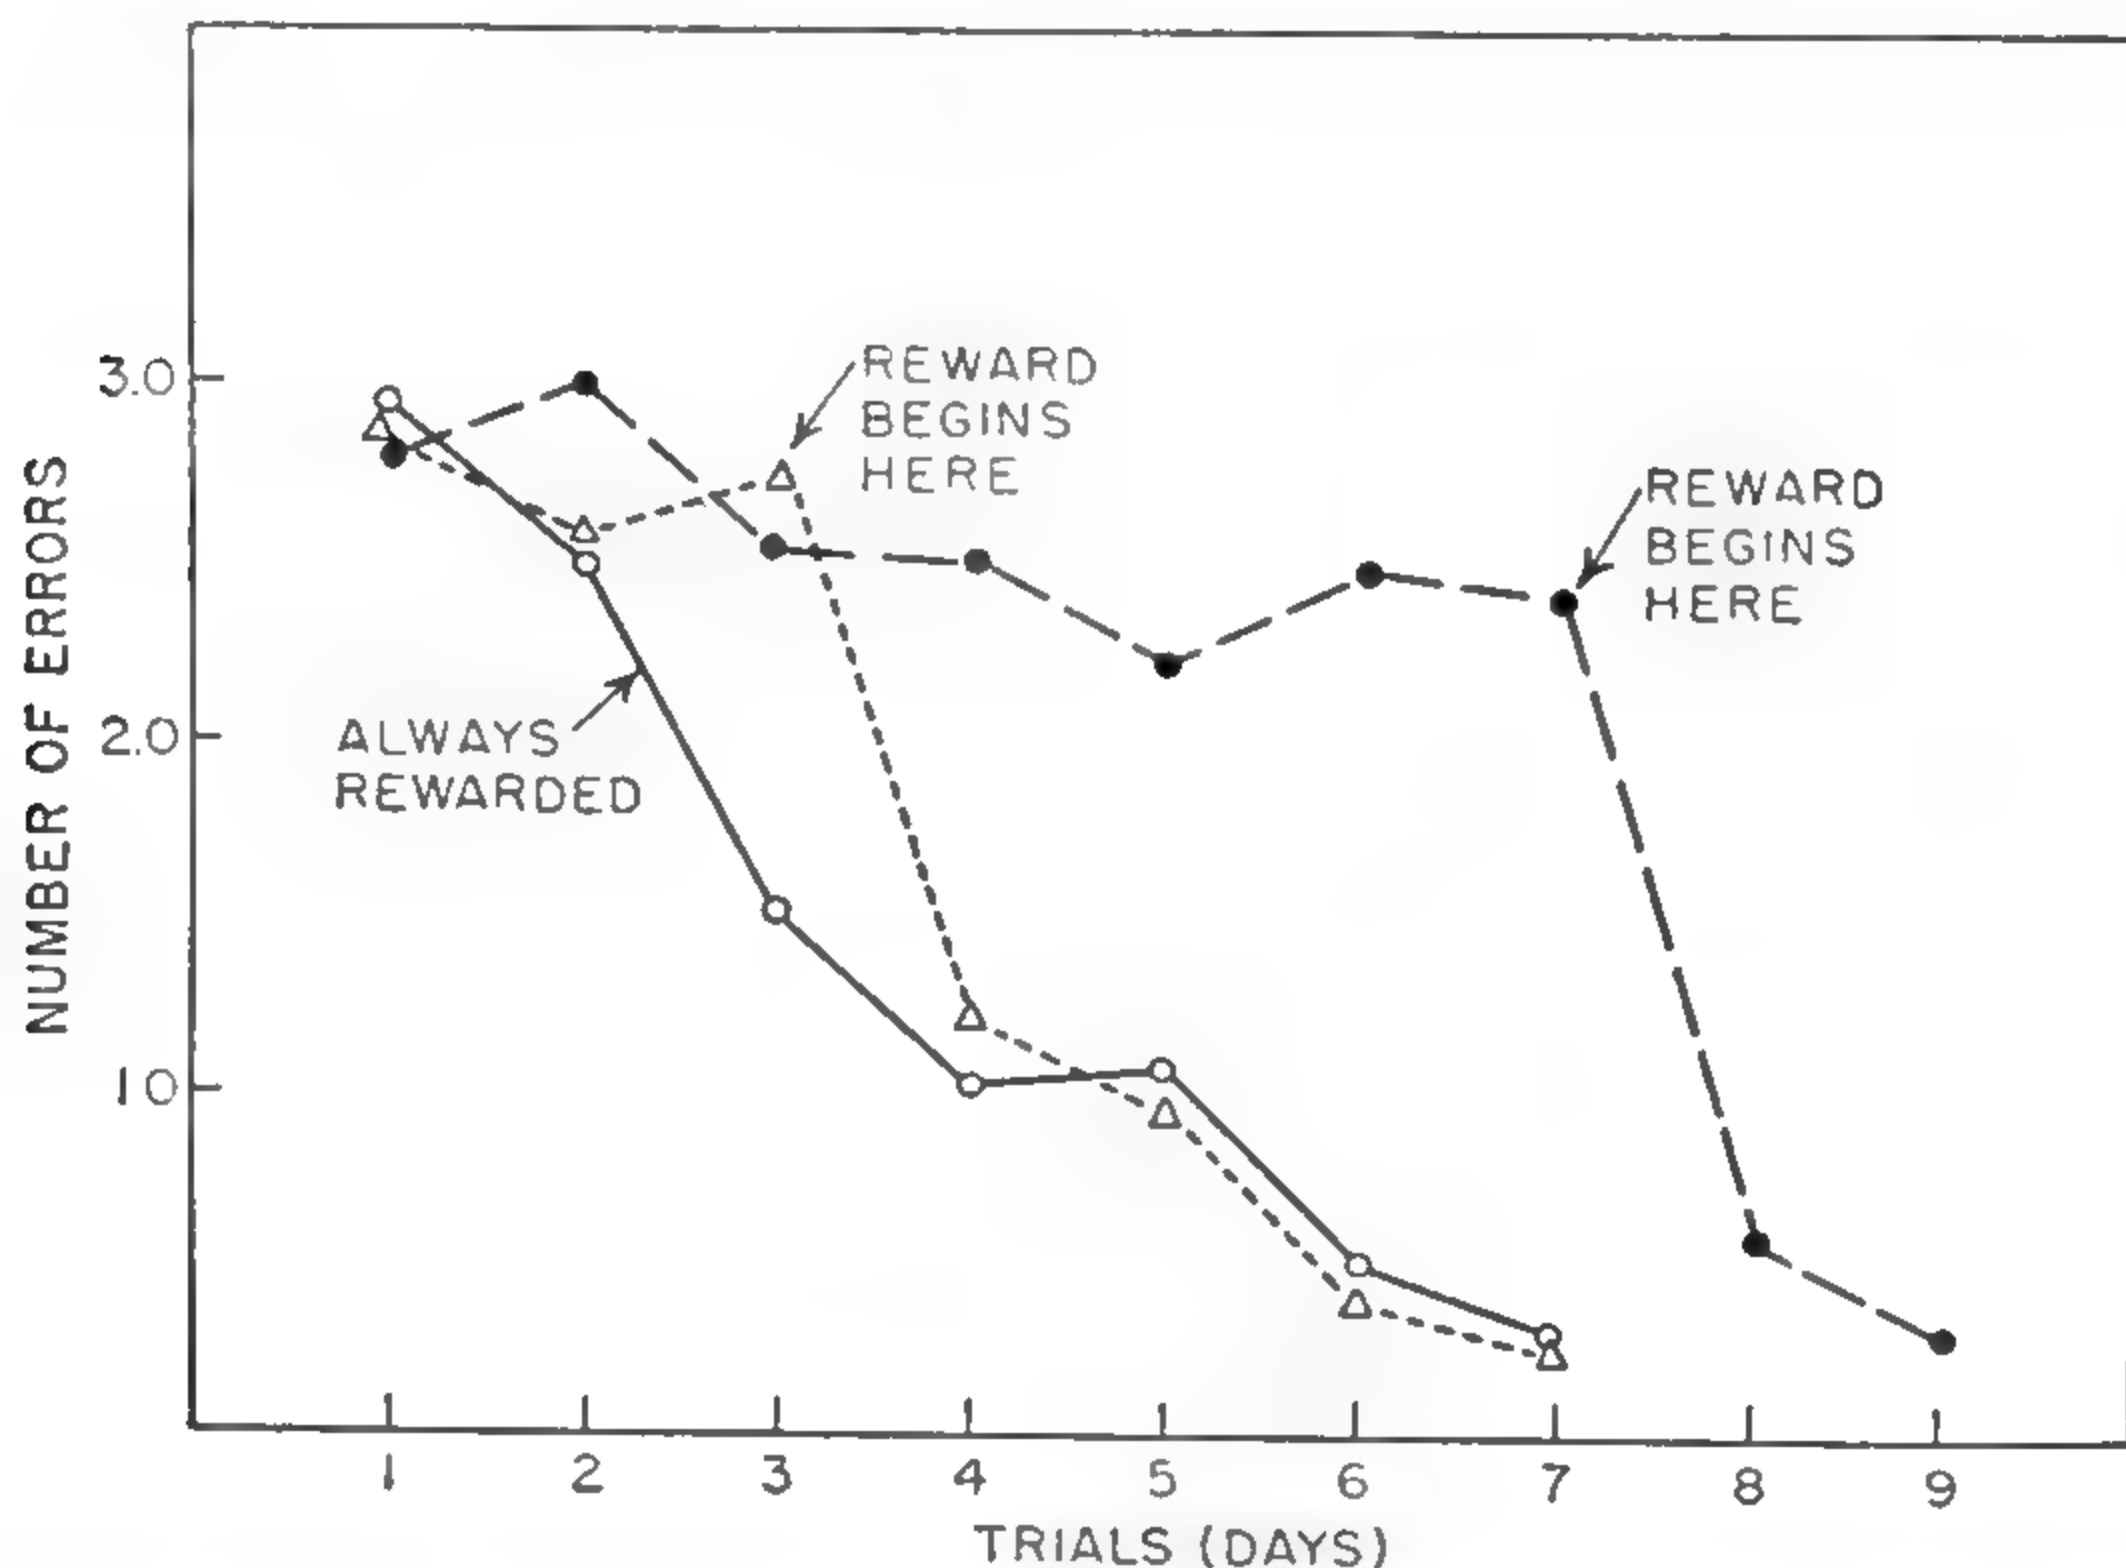


Figure 5. Number of errors during reinforced and nonreinforced trials in maze learning. When reinforcement is introduced, errors drop to a level comparable to that of animals continuously reinforced. (Data from Blodgett, 1929.)

for example, showed that the absence of food in the goal box does not necessarily mean the absence of reinforcement, for in this repetition animals that were not reinforced with food nevertheless showed improvement during the "latent" period, although they showed less than animals reinforced with food. Actually, as shown in Figure 5, Blodgett's animals learned a little as they explored the maze for seven trials without food. A later experiment (MacCorquodale and Meehl, 1951) revealed some of the things that were reinforcing in the absence of food, such as escape from situations that were anxiety-provoking for the rats.

It is clear that Blodgett's experiment is rather like the studies on amount of reinforcement. The group that did not receive food was weakly reinforced

and did not improve very much until given a strong reinforcement; then improvement occurred on one trial. Thus, Blodgett's study suggests, as do the studies on amount of reinforcement, that reinforcement affects performance, but not rate of learning.

Two more experiments on latent learning deserve mention. In these some of the critical weaknesses of Blodgett's study are removed. In one (Buxton, 1940), rats were permitted to live in a maze for a period of several days. They were placed in the maze at different positions and removed from it at different positions. Consequently, the animals explored all parts of the maze to about the same extent, and they were not differentially rewarded by being removed from any particular part of the maze.

The important result of this experiment is that it was not necessary to give the rats one rewarded trial before the sudden improvement in performance occurred. Instead of allowing the rats to run through the maze to discover the food the first time, as Blodgett had done, the experimenter put them directly into the goal box, so that they had a chance to see the food. When they were put at the beginning of the maze, most of the rats were able to find their way to the goal box without making many errors.

A second experiment (Seward, 1949) confirms this result, and these two experiments taken together show that it is possible for rats to learn the general plan of a simple maze when they are not differentially reinforced in any particular part of the maze. Though again, until the food is introduced as a reinforcement, performance does not show what the rat has learned.

Time of Reinforcement ✓

Another important variable in the instrumental learning of operant behavior is the time between the response and the reinforcement. This is not important in classical conditioning because there is no well-defined way of varying the time between the conditioned response and the reinforcement, though we may vary the time between the conditioned *stimulus* and the reinforcement (unconditioned stimulus). We shall examine this problem later. For the present we shall restrict ourselves to examples of operant behavior.

In an experiment by Perin (1943) we can see the basic nature of the relationship between time of reinforcement and strength of response. Perin used a Skinner box and rats. He slightly modified the usual procedure in the Skinner box so that a movement of the lever either to the right or left would produce a pellet of food. After the animals had learned to move

the rod in both directions and had displayed their preference for one direction, the apparatus was changed so that a movement of the rod in the preferred direction would give no food. A movement of the rod in the other direction, however, would cause the food pellet to be delivered.

The most important factor in this experiment is that the time between response and reinforcement was varied, ranging between 0 and 30 seconds for different rats. As illustrated in Figure 6, the results were clear. The sooner the reinforcement was delivered, the higher the rate of learning. As

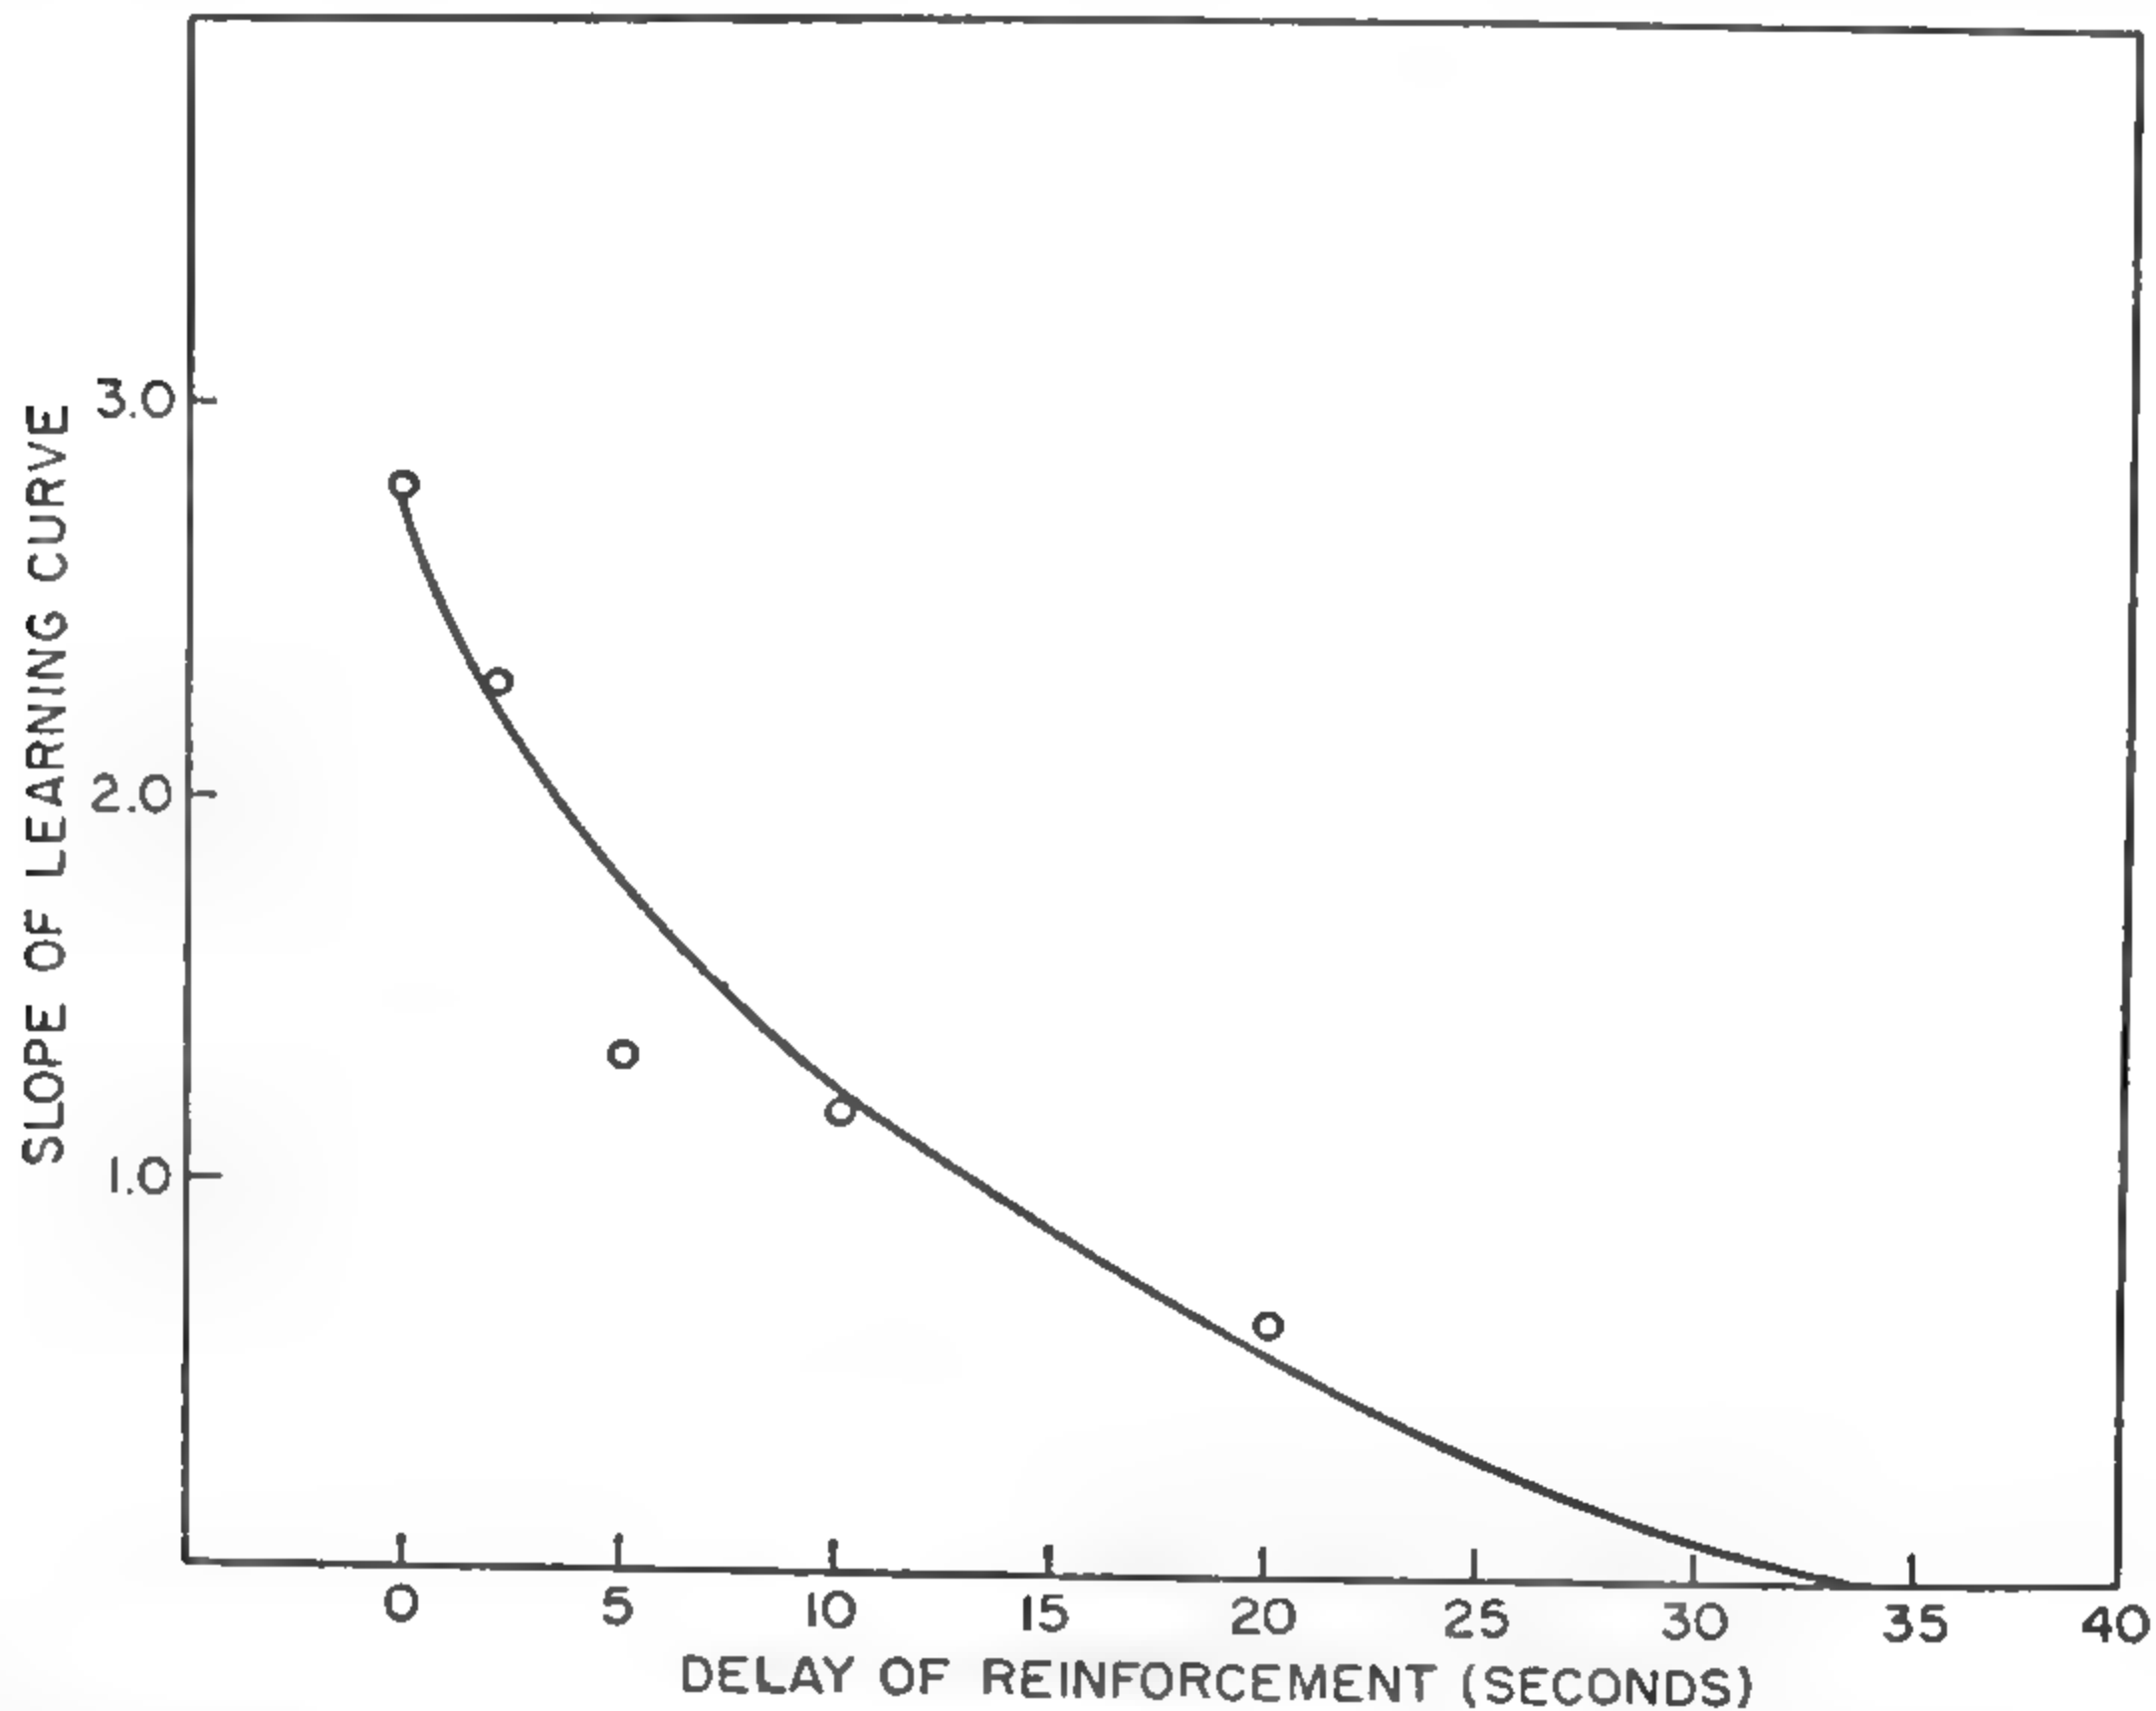


Figure 6. Slope of the learning curve as a function of time between response and reinforcement. The curve is extrapolated, on the basis of the five plotted points, to show zero learning at a delay between 30 and 45 seconds. (Data from Perin, 1942.)

a matter of fact, the rats with longer time delays learned so slowly that the experimenter predicted they would not learn at all if the time between response and the reinforcement were extended to about 45 seconds.

While Perin's experiment is a simple illustration of the basic relationship between time of reinforcement and response strength, it raises important questions. For one thing, there is a marked difference in the length of gradient achieved by Perin and that found by other investigators (see Hull, 1943). The results of the Perin study lead us to predict that if the delay had been 45 seconds, the rats would not have learned to press the lever at all. Other studies, however, show that under proper conditions the

delay between a response and a reinforcement can be extended to approximately 20 minutes and rats will still learn (Wolfe, 1934). We shall deal with this problem shortly.

The other problem created by this experiment is indicated by the label "Slope of learning curve," on the ordinate of the graph in Figure 6. Now this plotting suggests that time of reinforcement, rather than amount of reinforcement, affects both the limit and the rate of learning, and, indeed, this is the conclusion of an important theorist (Hull, 1951). Let us look at the implications of this finding.

Time of reinforcement and learning versus performance. Two experiments subsequent to Perin's show that the situation is complicated. Logan (1952) tried an approach similar to that used in the learning versus performance problem pertaining to amount of reinforcement; he trained rats to a steady level of performance under two different delays of reinforcement and then switched these delays. Initially, performance was poorer for a long delay than a short one. Animals that had been trained on a long delay gradually improved in performance when switched to a short delay, while the converse was true of those switched in the opposite direction. The fact that the change was gradual and not immediate suggests strongly that delay of reinforcement affects *both* learning and performance. A later experiment by Seward and Weldon (1953) shows essentially the same result.

These findings pose a problem since most theories find it difficult to accept the fact that *amount* of reinforcement affects only performance while *delay* of reinforcement affects both learning and performance. No simple explanation of this is possible, but let us look at a somewhat rudimentary version of one possible answer (Seward and Weldon, 1953).

One of the things animals have to learn in psychological laboratory tests is the relationship between what they do and the presence of the reinforcement. Put anthropomorphically, sometimes it seems as if the animals have a different notion about what they are supposed to be doing than the experimenters do. Skinner (1948) has pointed out that animals sometimes perform responses which are not necessary to obtain the reinforcement; this Skinner called "superstitious" behavior in animals. For example, a rat may appear to have the idea that it is necessary to turn around once counter-clockwise before pressing the lever. From Skinner's demonstration it appears likely that these false correlations between behavior and the reinforcement occur most frequently when reinforcement is delayed. In other words there is more uncertainty about which response is being reinforced when there is a delay between a particular response (such as lever pressing) and

the reinforcement. Since the reinforcement is sometimes connected by the animal to incidental behavior, the learned dependency between reinforcement and the response that actually produces it is weaker.

One implication of this explanation is that we should reconsider the relationship between amount of reinforcement and learning. It is all right to say that any amount of reinforcement will have about the same effect on learning as any other amount, but any response the animal learns will probably be weaker without reinforcement than it would be if it were reinforced. To put it anthropomorphically again, an animal might learn many things because they have equally interesting consequences, but because the animal learns these things more or less simultaneously, no one of them will be very strong. While it is clear that animals can learn in the absence of a consistent reinforcement, the learning may be weaker than it would be if such reinforcement were present. This seems to have been the basis for Tolman's statement that reinforcers are necessary to performance but serve only as "emphasizers" for learning (Tolman, 1938). If the situation is simple, it is likely that there will be no difference between reinforced and unreinforced learning, but in more complicated situations reinforcement may play a role in determining what animals learn.

Let us now turn to the other principal problem inherent in delay of reinforcement. Under some conditions reinforcements are effective after a long delay, while under others they lose their efficacy after a brief delay. Hull (1943) showed that the answer to this problem involved an important principle in the psychology of learning, the principle of *secondary reinforcement*.

Secondary reinforcement ✓ Secondary reinforcement applies to both classical and instrumental conditioning, so let us look at examples of secondary reinforcement in both types of learning.

In classical conditioning we can draw on an experiment of Pavlov's (1927). Here Pavlov used the ticking sound of a metronome as the conditioned stimulus and food as the unconditioned stimulus for salivation. After a few pairings of the metronome with the food, the salivary response became conditioned to the metronome. So much, of course, is only a typical conditioning experiment. Pavlov, however, carried this experiment one step further. He presented a black cardboard square to the dog. As we might expect, no saliva flowed to this stimulus. The black cardboard square was then paired with the ticking metronome ten times. By the tenth trial, the dog gave a salivary response to the black square, *even though it had never been paired with food, the unconditioned stimulus*. This response,

learned without the original unconditioned stimulus, Pavlov called a *higher-order* conditioned response, but today we generally speak of this as *secondary reinforcement*.

✓ An experiment on maze learning provides an example of secondary reinforcement in instrumental behavior (Saltzman, 1949). In this experiment rats were trained to run down a straight runway to a goal box that contained food. The goal box was either black or white. In one condition, the goal box always contained food when it was black but did not when it was white. After this training, the rats were taught a simple maze in which they had to choose between two pathways, one leading to a black goal box and the other to a white one. The rats learned to go to the black goal box, *even though it was never baited with food*. Thus it is clear that the black goal box itself was reinforcing because it had been paired with a primary reinforcement, food, in another situation. The black goal box was thus a secondary reinforcement.

A general definition of secondary reinforcement is that it is any stimulus that acquires reinforcing power because it is paired with a primary or unlearned reinforcement. It is easy to see that the concept of secondary reinforcement is important, for it accounts for learning in which there seems to be no reinforcement based upon a known biological drive. However, we may ask, what does secondary reinforcement have to do with the time between a response and its reinforcement? Since this is an important question, let us examine it in detail.

Time of reinforcement and secondary reinforcement. The principle of secondary reinforcement enables us to describe why delaying reinforcement only sometimes has a deleterious effect upon learning. We mentioned earlier an experiment on lever pressing in which learning practically disappeared if the reinforcement was delayed much more than half a minute, while in another experiment on maze learning the reinforcement could be delayed for almost half an hour and the rats would still learn. Let us compare the two.

In the maze learning experiment (Wolfe, 1934) the animals had to learn to go to one arm of a T-shaped maze. The goal box at the end of this arm contained food, and the identical goal box in the other arm did not. The maze had doors just before the goal boxes, so that the animals could be retained in the arms of the T before they were allowed to enter the goal box. Thus the arms of the T served as retention chambers.

It is the retention chamber that explains why delay did so little harm to learning in this situation. Since the retention chamber in one arm always

immediately preceded food, it came to be a specific secondary reinforcement. Thus, no matter how long the animals were delayed before they entered the goal box, they could learn to turn to the side with food—simply because of the secondary reinforcement associated with this particular retention chamber.

In the lever-pressing experiment (Perin, 1943), the experimenter went to a good deal of trouble to eliminate all secondary reinforcement. He built the apparatus so that the rod was removed immediately after the rat had pressed it; this stimulus, which might have become a secondary reinforcement, was therefore never present when the rat actually got the food (except when the delay between response and reinforcement was zero). Thus there was minimal opportunity for the lever to become a secondary reinforcement. As a result, the amount of delay between the primary reinforcement and the response had to be short if the animal was to learn, for there was little or no secondary reinforcement to bridge the gap.

The gradient of reinforcement. The relation between response strength and delay of reinforcement has been called the *gradient of reinforcement*. This refers to the fact that the more immediate the reinforcement, the stronger the response. We have seen that if no opportunity for secondary reinforcement is present, and if there is any delay, however short, between response and reinforcement, the animal will not show that it has learned.

Spence (1947) has made the suggestion that any delay between response and reinforcement is mediated by secondary reinforcement. That is to say, when all cues for secondary reinforcement are eliminated, there may be no learning unless the response and the primary reinforcement occur simultaneously. This is probably something of an exaggeration. We have already seen that reinforcement seems to exert a primary influence on performance and only a secondary one on learning. So it is that animals will sometimes learn with no apparent reinforcement, though they will not perform well. Nevertheless, Spence's point is well taken. Most of the delay between a reinforcement and a particular response is probably bridged by secondary reinforcement. A variety of experiments have shown that it is possible to change the extent of the gradient of reinforcement by changing the possible secondary reinforcements (Perkins, 1947; Grice, 1948; Smith, 1951).

THEORIES OF REINFORCEMENT

We have had little to say thus far about theories of reinforcement, and yet for many years these have been of major concern to psychologists inter-

ested in learning. Although it is certainly true that the question of why reinforcement occurs is important, we have deferred discussion of this problem because many basic issues are still unsettled. On several important matters we cannot even state any very firm conclusions. Some questions, however, have been settled by experimental work, and theorists have come to some agreement on still others. Let us look at the basic issues and see how some of them have been tentatively settled, while others have remained controversial.

Basic Issues in Theories of Reinforcement

In the following section we shall examine some of the more important issues involved in the various theories of reinforcement and the attempts that have been made to resolve them. Two problems are of paramount importance. One is the question of why certain stimuli appear to be reinforcing and others do not. The other is whether reinforcement determines learning or performance. In addition there is a subsidiary problem of whether or not the same principles of reinforcement apply to both respondent and operant behavior.

Why are some stimuli reinforcing? It is fairly clear that there are only a limited number of stimuli that are reinforcing to every animal. Animal trainers know this, of course, and they are careful to reward the behavior they want to teach their animals with the proper reinforcement. If we are teaching tricks to dogs, a simple pat on the head will reinforce most of the time, though food is probably better. With the domestic cat, a pat on the head is usually not enough. Indeed, it is possible that the difficulty people have in teaching cats may be due in part to a more limited class of reinforcing stimuli.

Although it does not sound very enlightening to say that reinforcing stimuli are things that animals like, this common-sense definition contains the germ of some of our most important theoretical ideas about reinforcement.

The origins of present interpretations of why some stimuli reinforce animals and others do not probably come from the theory of natural selection. Darwin and his successors tried to explain the process of species evolution in terms of natural selective breeding. The adaptive characteristics of organisms survive, said Darwin, and the unadaptive ones die out. Naturalists have pointed out many examples of anatomical and physiological characteristics that are adaptive and that seem to be perpetuated through selective breeding. Thus in evolution, biological changes in organisms have "good"

and “bad” effects, and, according to the natural selection theory, only the changes with “good” effects survive.

It was apparent to many early workers in evolution that behavioral patterns also provide good examples of the survival of adaptive mechanisms. Furthermore, in the lifetime of an individual there appears to be a process of selective adaptation. Organisms seem to learn those things which are useful—they learn the location of food supplies, water, hiding places, etc. The early literature of comparative psychology abounds with anecdotes illustrating the survival value of the things animals learn.

Thus, the argument goes, adaptive behavior (that which preserves the animal from harm and keeps it fed and sheltered) survives and becomes learned, whereas unadaptive behavior does not. This concept was combined with the hedonistic ideas of Utilitarianism. Pleasure, it was held, is associated with adaptive behavior and pain with nonadaptive behavior. Thus pleasure and pain were brought into an adaptive theory.

This kind of notion—a behavioral hedonism combined with the survival of adaptive traits—has had an enormous influence on the theory of learning and behavior. It is hard to find theorists who deny that reinforcing stimuli are connected with adaptive mechanisms in the animal. Many theorists deny the universal importance of adaptive reinforcement, however, and others say that, while it determines what the animal will do, it has nothing to do with what the animal learns. Thus we face once more this issue of learning versus performance.

Reinforcement in learning and performance. The question of whether reinforcement determines learning or merely performance is a relatively new one. Most of the older investigators did not make a distinction. E. L. Thorndike, one of the first experimental psychologists to explore the adaptive characteristics of reinforced behavior, never clearly stated whether he thought that reinforcement affected performance or both learning and performance. In his *Principle of Effect* (1898), which led to our current operational notions of reinforcement, he said that rewards tend to *stamp in* behavior. This and other of his statements imply that reinforcement acts upon both what the animal learns and what it does—indeed that reinforcement is necessary for learning.

This implication from Thorndike had considerable influence, and for a long time the dominant view was that an adaptive reinforcement was necessary for both learning and performance. Hull (1943) was the most influential theorist to adopt this viewpoint. He held that adaptive reinforcement

was necessary for the learning of both operant and respondent behavior. Hull was a quantitative-learning theorist and wrote nearly all his hypotheses as equations. It is interesting to note that near the end of his career Hull drastically revised his theory by rewriting his equations so that performance but not learning was determined by reinforcement (Hull, 1951).

The need-reduction theory of reinforcement. Hull's original belief in the importance of reinforcement for both learning and performance led him to speculate on the mechanism of reinforcement, and his conjectures have provided a specific theory that is valuable whether reinforcement affects learning and performance or just performance. This is the *need-reduction* theory of reinforcement.

The need-reduction theory states that reinforcements are always related to specific motive states in animals and men. They reduce the level of organic or secondary need derived by a process of learning from some organic need. Thus food is reinforcing for a hungry animal, just as water is for a thirsty one. Other drives, sex, for example, are also reduced by appropriate reinforcement. Most experimental work, however, tends to deal with hunger and thirst since both the level of drive and its appropriate reinforcement are easier to control.

There is little doubt that drives and their reduction control the strength of responses. Whether all behavior is directly or indirectly controlled by need reduction is impossible to answer, but there is experimental evidence of the control exerted by factors such as reduction of hunger. An experiment by Miller and Kessen (1952), for example, shows that milk introduced *directly* into the stomach of rats by way of a fistula can serve as a reward in maze learning. The rats in this experiment could not taste the milk, so it was obvious that the milk was not reinforcing because of its taste or because of "consummatory activity"; furthermore, a control experiment in which saline was used instead of milk made it clear that the critical factor is not just something in the stomach but rather the nutritive value of food.

Not all behavior is controlled by primary reinforcement, according to this theory. For, of course, secondary reinforcements can be built upon primary reinforcement. Secondary reinforcements, moreover, involve secondary drives (Miller and Dollard, 1941). It is uncertain whether or not secondary reinforcements reduce secondary drives as primary reinforcements do primary drives. In the case of fear (a secondary drive) it is thought that escape from fear (a secondary reinforcement) does result in the reduction of the secondary drive, but this seems to be a somewhat special case, so we

cannot make any general statement from it. It is quite possible, however, that secondary reinforcements do involve some reduction in a need analogous to the reduction of a primary drive by a primary reinforcement.

In recent years investigators have found evidence for a particular system in the brain that, when stimulated, reinforces behavior (Olds and Milner, 1954). Since such stimulation does not reduce a biological need, it is argued that need reduction itself is not necessary for reinforcement. In the usual course of things, however, need reduction may activate this reward system in the brain (Olds, 1955). Although need reduction is what normally dictates animal behavior, it is not the essential element. Perhaps activation of a particular part of the brain is the important factor.

Furthermore, there is reason to believe that there is a direct sensory, or hedonic, kind of reinforcement (Young and Shuford, 1954). That is to say, some things are reinforcing because they have the right kind of sensory consequences, not because they satisfy some biological need. Thus saccharin is reinforcing for rats (Sheffield and Roby, 1950) even though it has no nutritive value. Presumably it is reinforcing only because it affects the right taste receptors.

Thus, while the original need-reduction theory seems rather naïve now, it is only because we have a better understanding of the possible ways in which reinforcement controls behavior. Need-reducing reinforcements affect behavior to some extent, but they are only part of the story. Furthermore, along with Hull, most theorists have come to the view that need reduction and other possible sources of reinforcement, such as sensory effect, control what animals do but not what they learn. They may affect complex learning in one way or another, but reinforcing mechanisms are not essential to the learning process.

An association theory of learning. We must then ask, Why does learning take place? The answer is very old; in its simplest form it extends back to Aristotle. We can say that the only condition that seems basic to learning is that two things become associated. To make this statement intelligible, however, we have to state what kind of things become associated and under what conditions. On this point, there are several different views.

In contemporary psychological theory one of the most common ideas is that stimuli and responses are the associated factors. Stimuli come to elicit responses that they previously had not elicited because, perhaps quite by accident, these stimuli occur simultaneously with certain responses. The model for this process is the classical conditioned response. Here we elicit an unconditioned response by using an appropriate unconditioned stimulus at

the same time that we present a conditioned stimulus. Because the conditioned stimulus is simultaneous with the response, it comes, through association, to elicit that response.

This association theory can be applied directly to classical conditioned responses. It also applies to instrumental conditioned responses, even though we cannot always produce experimentally stimuli that are simultaneous with the response to be conditioned. Suppose, for example, we are trying to teach a dog to retrieve something when we say, "Go get it, boy." We can get dogs to retrieve easily enough, but it is much more difficult to get them to retrieve only on a given signal. Let us see how this kind of association theory says the process occurs.

The best-known advocate of this kind of association theory has been E. R. Guthrie, who has this to say about association: "A combination of stimuli which has accompanied a movement will on its recurrence tend to be followed by that movement" (Guthrie, 1935, 1952). For Guthrie the classical conditioned response is the model for all learning processes. In most cases, however, we cannot observe conditioned stimuli (for example, the rat pressing the bar in the Skinner box); and so in Guthrie's theory they take on the character of hypothetical constructs. That is to say, we invent some conditioned stimuli for a particular set of learned responses. In the case of lever pressing, the conditioned stimuli are kinesthetic, according to Guthrie (1952). In other words, the rat makes a series of movements, these stimulate the sense organs of the muscles, and this combination of muscular sensations serves as a conditioned stimulus.

Thus, according to this particular version of association theory, if we wish to teach someone an instrumental response, we try to arrange it so that the stimulus and the response occur together. Obviously with such a principle of learning the distinction between classical and instrumental conditioned responses becomes purely a matter of convenience. Guthrie (1952) believes that this distinction may even be a bit pernicious, since it tends to make us accept instrumental behavior as purely "spontaneous" and prevents us from looking for stimuli which might cause the behavior.

Strangely enough, however, Guthrie's major experimental study is on instrumental learning. Guthrie is dissatisfied with the typical experiment on instrumental behavior, because, he claims, it examines the *results* of behavior, not the behavior itself. Take the lever-pressing case, for example. In the Skinner box we find out how many times per minute the rat presses the lever; ordinarily we are not interested in the *way* in which the lever is manipulated by the animal. Guthrie objects to this approach, for he believes

that it is the association of particular *movements* of the animal with stimuli that is really important and that the result of the movement is only secondary. Consequently in an experimental situation, Guthrie and Horton (1946) studied the way in which cats learned to escape from a puzzle box. Instead of merely recording the time it took the cats to get out of the box or the number of times they escaped, Guthrie and Horton observed the cats in action and took motion pictures. One of the things that impressed them was the degree of behavioral stereotypy. In many cases the cats' movements were almost identical from trial to trial. This indicated to Guthrie that there was a close association between the cues and the movements involved.

One of the implications of Guthrie's notion of association of stimulus and response through contiguous exposure is that it takes only one pairing for conditioning to occur. This seems to go against everything we have said thus far, since curves show that learning is a gradual process. Guthrie has a good answer for this: An ordinary instrumental act (such as lever pressing or escaping from a puzzle box) involves many different movements. Thus it may take many trials for the animal to perfect a *pattern* of movements that is eventually successful in performing the act. Most behavior involves the association of many stimulus-response units. We are constantly learning and unlearning these associations, so that there is a continuous fluctuation in our behavior. Therefore it takes time to perfect a series of movements that accomplish a particular result.

This last point of Guthrie's has been elaborated into an important mathematical theory of learning by Estes and Burke (Estes, 1950; Estes and Burke, 1953). In this theory the behavioral variability from trial to trial is attributed to the effective stimulus. On each trial the animal is stimulated with a *sample* of the available stimuli. The samples are partially overlapping on each successive trial, but the response that occurs on any one trial is that which is conditioned to 100 per cent of the sample stimuli on the particular trial. The fact that there will be sampling variability from trial to trial in the effective stimuli means that new unconditioned elements as well as old conditioned ones will be present each time.

For example, all the stimuli in the Skinner box constitute a population of stimuli; at any one moment, however, the rat is stimulated only by a sample of these. Sometimes these samples will overlap a great deal, and occasionally not at all. From the laws of probability one can generate a learning curve for lever pressing based on the frequency with which successive stimulus samples will contain common elements.

We shall have occasion to examine some of the mathematical properties of this theory of Estes and Burke later. At present we should note that it suggests, as does its parent theory by Guthrie, that all conditioning is essentially classical and that the distinction between classical and operant conditioning is only an experimental convenience, not a basic difference. Both of these theories assign all conditioning to the classical variety because they suggest that the basis of learning is the establishment of a correlation between some novel stimulus and a response it has not hitherto evoked. By confirming this new stimulus-response connection enough times we can make sure that the elements of the stimulus population and the components of the response are sufficiently associated to produce efficient performance. We must refer to this association as a *correlation* of stimulus and response, because in instrumental behavior we cannot always cause the association to be present; we are only capable of devising a situation in which there will be a correlation between a stimulus situation and a "spontaneous" response. Thus, this particular variety of association theory says that all learning is exemplified by classical conditioning, and only our lack of control over unconditioned stimuli for instrumental or operant behavior prevents us from doing experiments on instrumental behavior that are exactly parallel to those on respondent behavior.

In this kind of association theory, therefore, the function of reinforcement is simply the mechanical one of providing an end to the stimulus-response sequence. Reward serves to prevent the animal from unlearning what it has already learned by keeping it from reacting in other ways to the stimuli that lead to the response. Food at the end of a maze keeps the animal there instead of allowing it to wander through the maze unlearning proper turns and learning new turns to some blind alley. Thus, the rat eventually does learn the maze. Reward only serves to keep the rat from unlearning what it has learned. Incidentally, this suggests that the rat would learn the maze if we simply removed it from the goal box at the end of the trial instead of giving it food. Seward (1942) tested this notion with the Skinner box—that is, he simply removed the animals from the box after they pressed the lever—and they did learn, though not so well as food-rewarded animals.

It is difficult to see how the stimulus-response association theory could predict performance in latent-learning cases without some additional principles, for latent-learning experiments show that animals learn more than they perform during the latent phase. This particular association theory, as we have seen, puts its greatest emphasis upon performance—the movements made by the animal.

Another association theory. The topic of latent learning brings us to another association theory, one that is radically different. As a matter of fact, it is probably unfair to call it an association theory at all. It has one thing in common with the association theory that we have been talking about, though. The idea is that need-reducing reinforcement has little to do with learning, but that learning is a matter of contiguity. In this case, however, the contiguity is not necessarily between stimulus and response.

The foremost proponent of this view, that is, of an associative theory that does not use the classical conditioned response as its model, is E. C. Tolman. In a series of publications,² he has developed and expanded what he calls a "purposive" theory of learning. This theory emphasizes "sign learning" as opposed to response learning. This means that animals usually learn (when running a maze, for example) the succession of stimuli that lead to the goal. In one of Tolman's characteristically epigrammatical phrases, the rats make "cognitive maps" of the maze (1948). Because animals learn *relationships* between stimuli rather than *responses*, he was able to predict the sudden shifts in performance obtained in complicated latent-learning experiments.

Although there are many difficulties in Tolman's theory, it has led to some interesting experiments. For the present we shall be content to mention it merely as a theory of learning that does not emphasize the need for reinforcement to "stamp in" learned responses. Reinforcement may help learned responses but learning itself can take place without it.

Perhaps the most important contribution of Tolman has been his distinction between learning and performance. He made this distinction early, and since then, perhaps because of the latent-learning experiments, most theorists have adopted it. Hull, for example, accepted it in his book written in 1943, though he still defended the position that reinforcement was necessary for learning. In his last publications (1951, 1952), however, Hull accepted Tolman's distinction in principle by admitting that learning was a function of neither reinforcement nor drive.

The current status of theories of learning and reinforcement. There is no doubt that we can characterize most instrumental behavior as reinforcement dependent. If we look at most examples of laboratory behavior, we see that there is usually some special event necessary to produce and maintain the behavior in strength. This can be giving a pellet of food to a hungry rat or allowing a monkey to glimpse a toy train in operation (Butler, 1954).

² Many of Tolman's papers have been brought together in a collection (see Tolman, 1951).

One of the arts in the control of behavior is knowing the range of reinforcements possible for a given animal and how to administer them for optimal results. Classical conditioning, of course, depends upon the reinforcing property of the unconditioned stimulus. As does instrumental behavior, a classical conditioned response tends to disappear if it is not occasionally reinforced with an unconditioned stimulus (see Chapter 3). There is no doubt about the importance of the reinforcement operation in the production and control of behavior.

There is now strong evidence, however, that reinforcement is not necessary for learning, although it does probably modify learning. Thus reinforced learning in a complicated problem, such as learning a long maze, is likely to be different from unreinforced learning. This is not because reinforcement is necessary to learning, but because reinforcement indirectly modifies what animals learn by determining what they do.

Thus, although we described earlier how learning increases as a function of *number* of reinforcements, we should now modify this to see how learning increases as a function of the number of exposures or experiences in the learning situation. Usually we reinforce successive tries, so that an increase in the number of reinforcements is correlated with the change in behavior. The essential requirement for a series of changes in behavior is simply that the animal behave in a certain way on successive exposures to the situation. All that is necessary to maze learning is that the animal run to the goal box a certain number of times. We can usually accomplish this most efficiently by reinforcing the animal. Thus the important variable in learning is the number of trials which are given an animal on a particular problem.

The distinction between classical and instrumental conditioning. We have seen in Guthrie's stimulus-response association theory that the distinction between classical and instrumental conditioning is a matter of convenience. Other theorists, however, maintain that there is a fundamental difference between them. Mowrer (1947), for example, points to the fact that the reinforcement in classical conditioning is not need-reducing, while such reinforcement is usually necessary to produce instrumental or operant behavior. Other investigators have pointed out that classical conditioning is characteristic of simple reflexes, and instrumental conditioning of voluntary behavior. Certainly, at the present time, we cannot come to a definite conclusion as to whether the distinction between classical and instrumental conditioning is real or merely convenient. This problem will come up in later chapters, and we shall see that it is of great current interest.

TEMPORAL ASSOCIATIVE FACTORS IN CONDITIONING

Since reinforcement is not essential to learning, the most fundamental conditions we can describe are associations—associations between stimulus and stimulus, stimulus and response, and response and response. Association means that two events, stimuli or responses, must occur close together in time. We have already examined the data on response-reinforcement associations, but not the stimulus-stimulus relationships in classical conditioning. In this section we shall look at temporal factors in classical conditioning. These are of basic importance because classical conditioning is the simplest example of associative learning.

Pavlov's studies on time relations in conditioning. Pavlov (1927) used two different techniques in the study of time relations in classical conditioning. In one, the conditioned and unconditioned stimuli overlapped each other in time. The conditioned stimulus came either before the unconditioned stimulus or the two were simultaneous, but the essential condition was that they overlapped. In the other technique, the two stimuli did not overlap. Usually the conditioned stimulus was presented first, and then, after an interval, the unconditioned stimulus.

Pavlov thought that simultaneous presentation of the conditioned and unconditioned stimuli produced the most rapid conditioning, and in all the experiments in his laboratory, they were presented in this way. Following this, he gradually lengthened the interval between the conditioned and unconditioned stimuli, or allowed longer exposures of the conditioned stimulus before he presented the unconditioned stimulus.

When the unconditioned stimulus was delayed, Pavlov found that animals could learn to delay the onset of the conditioned response for long periods of time. They would delay giving a conditioned response until just before the unconditioned stimulus. Pavlov took this to mean that the conditioned response was actively inhibited. He used the delayed conditioned response as an example of what he called "internal inhibition." He also found that animals could delay the onset of the conditioned response when the conditioned stimulus was presented and then withdrawn before the unconditioned stimulus came on. This Pavlov called the "trace" conditioned response, and he considered it also to be an example of internal inhibition. He found, moreover, that it was more difficult to establish the trace conditioned response than the delayed conditioned response.

These observations are important in the present context because they indicate associations between conditioned stimuli and conditioned responses to be extraordinarily flexible. Such associations do not have to be based on simultaneity or, indeed, upon close temporal association. Pavlov could introduce delays of many minutes in the trace conditioning procedure and still produce a conditioned response; the conditioned stimulus would come on briefly and then, after a wait, the conditioned response would occur. For the production of such delays it is apparently essential that nothing else occur between the conditioned stimulus and its response. Pavlov's experiments were performed in sound-deadened rooms in order to control the stimuli impinging on the animal, so that the animal would not be stimulated in any way between the occurrence of the conditioned stimulus and the time for the conditioned response. Such an absolutely controlled environment seems to be necessary for very long delays. The animals themselves came to inhibit activity in the time interval, and in some instances appeared to go to sleep. Thus, as appears to be the case with delay between responses and reinforcement in instrumental conditioning, if nothing interferes with the association the delay can be bridged.

Rate of learning and delay of the unconditioned stimulus. A number of investigators have studied the rate of conditioning as a function of the time interval between the conditioned and unconditioned stimuli. In these studies the highest rate of conditioning occurs when the onset of the conditioned stimulus just *precedes* that of the unconditioned stimulus (Spooner and Kellogg, 1947; Kimble, 1947; Moeller, 1954).³ These studies were all performed with human subjects; the conditioned responses were variously a hand movement, the blinking of the eyelid, and a galvanic skin response. Despite these differences, the studies all agree that the highest rate of conditioning occurs when the conditioned stimulus precedes the unconditioned stimulus by approximately half a second. Figure 7 shows the data on the conditioned hand movement, together with some data from earlier studies of the same response (Wolfe, 1930, 1932).

One curious case of the time relation between the conditioned and unconditioned stimuli is shown in Figure 7. This is *backward* conditioning. In backward conditioning the unconditioned stimulus precedes the conditioned stimulus. Pavlov thought that conditioning could not take place in this case; however, the data presented in Figure 7 show some evidence for backward

³ The experiment by Kimble (1947) is not strictly an example of classical conditioning, though it is sufficiently close to be considered so for present purposes.

conditioning. It is probably not true associative conditioning, however, but more likely a sensitization of the whole response system. The unconditioned stimulus (an electric shock in the cases in which backward conditioning is found) sensitizes the subject so that he will give a response to almost any stimulus. This interpretation agrees with the data of Spooner and Kellogg (1947). They find that the number of backward conditioned responses is high early in training but decreases during it. This is because the subjects readily adapt to stimuli other than the shock.

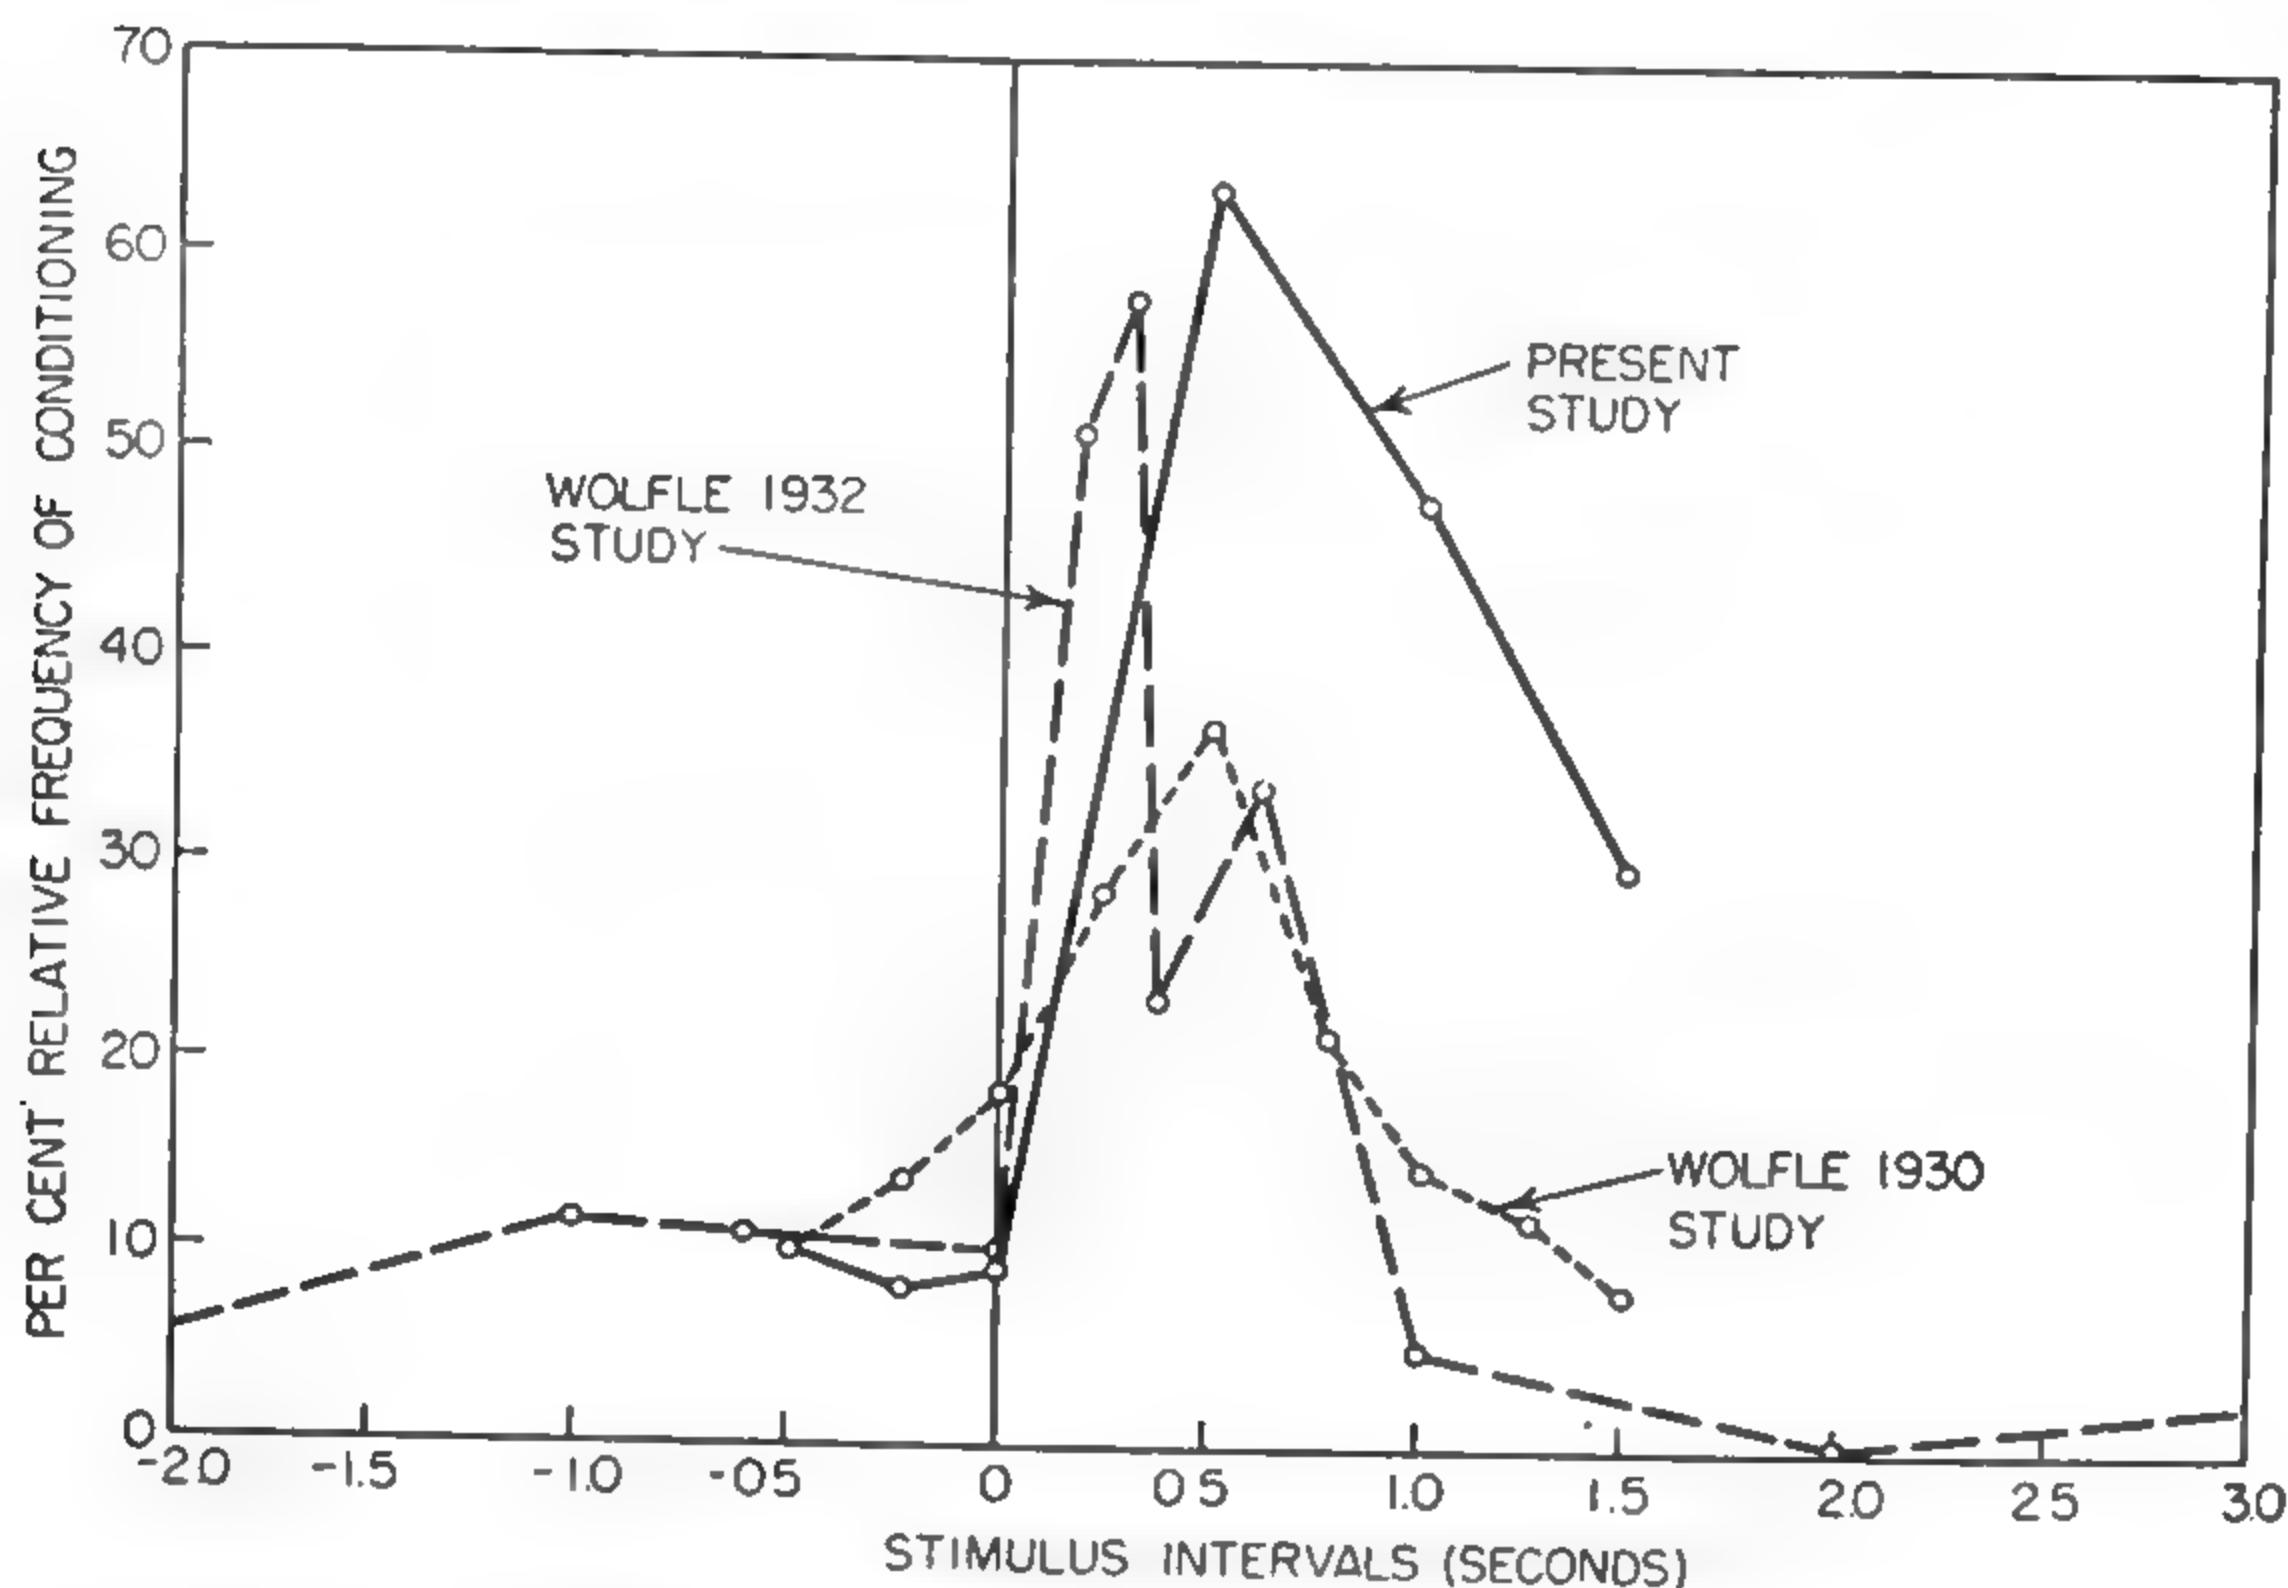


Figure 7. Results of the Spooner and Kellogg (1947) and the Wolfe (1930, 1931) studies on the time interval between conditioned and unconditioned stimuli. The curve labeled "present study" is from the data of Spooner and Kellogg. Notice that the maximum frequency of conditioned responses occurs when the conditioned stimulus precedes the unconditioned stimulus by a fraction of a second. Backward conditioned responses rarely occur.

Rate of conditioning and association. All studies show that conditioning is less efficient when the stimuli are presented together than when the conditioned stimulus precedes the unconditioned stimulus by a brief interval. This fact is fundamental to the interpretation of the associative processes. The association is not always between the stimuli, but may be between one stimulus and the *trace*, or aftereffects, of another. Just why asynchrony is best is not clearly understood at present. The brevity of the optimal time interval between the onset of the conditioned stimulus and unconditioned stimulus suggests that this effect has to do with the time of events in the central nervous system. The important fact, however, is that the rate of

conditioning drops rapidly as the time interval between the stimuli is increased to even a few seconds (Moeller, 1954). In trace conditioning this interval may be extended, but only through tedious training under conditions in which there is a minimum of external stimulation. Thus while the duration of the stimulus trace is probably long, it is easily obscured by other stimuli, and it is therefore difficult to obtain the kind of associations characteristic of classical conditioning except with very short time intervals.

CHAPTER 3

EXTINCTION AND PARTIAL REINFORCEMENT

When learned acts are no longer reinforced, they usually diminish in strength. Frequently, for example, colicky babies learn to cry in order to obtain extra attention. If this happens, parents are sometimes advised to avoid picking up the baby when it cries so that it will stop crying for attention. It is doubtful whether many parents last out the process, but the advice itself is testimony to the common knowledge of the effects of withdrawal of reinforcement.

The effects of removing reinforcement from a particular bit of behavior are quite complicated. The general effect—diminished strength of response caused by a failure of reinforcement—is called extinction, and in the first part of this chapter we shall look at the way in which this is modified by different variables. In later sections we shall examine the more general problem of the effects of intermittently given reinforcement—we shall see what happens when a given response is reinforced some, but not all, of the time.

EXTINCTION

In the simplest examples of learning, a response is reinforced every time it occurs. Thus, each time the rat in the Skinner box presses the lever, it receives a pellet of food. After the rat has learned about the correlation between food and pressing the lever, it will continue to press the lever at a steady rate as long as it remains hungry. If the reinforcement is discontinued, the rat will continue to respond for a while, but at a gradually diminishing rate and with increasing irregularity. Figure 8 shows how with-

drawal of reinforcement affects rate of responding in the Skinner box. From this figure it is clear that the rate of responding is high at the beginning (right after reinforcement is withdrawn) and that it gradually diminishes, so that by the end of the first hour without reinforcement, the rat is going long periods of time without giving a response. It is also clear that the rate of responding is quite irregular during extinction (compared with the highly regular rate under reinforcement shown in Figure 2). The rat may press the lever rapidly for a few minutes and then cease to press it altogether for a period of time.

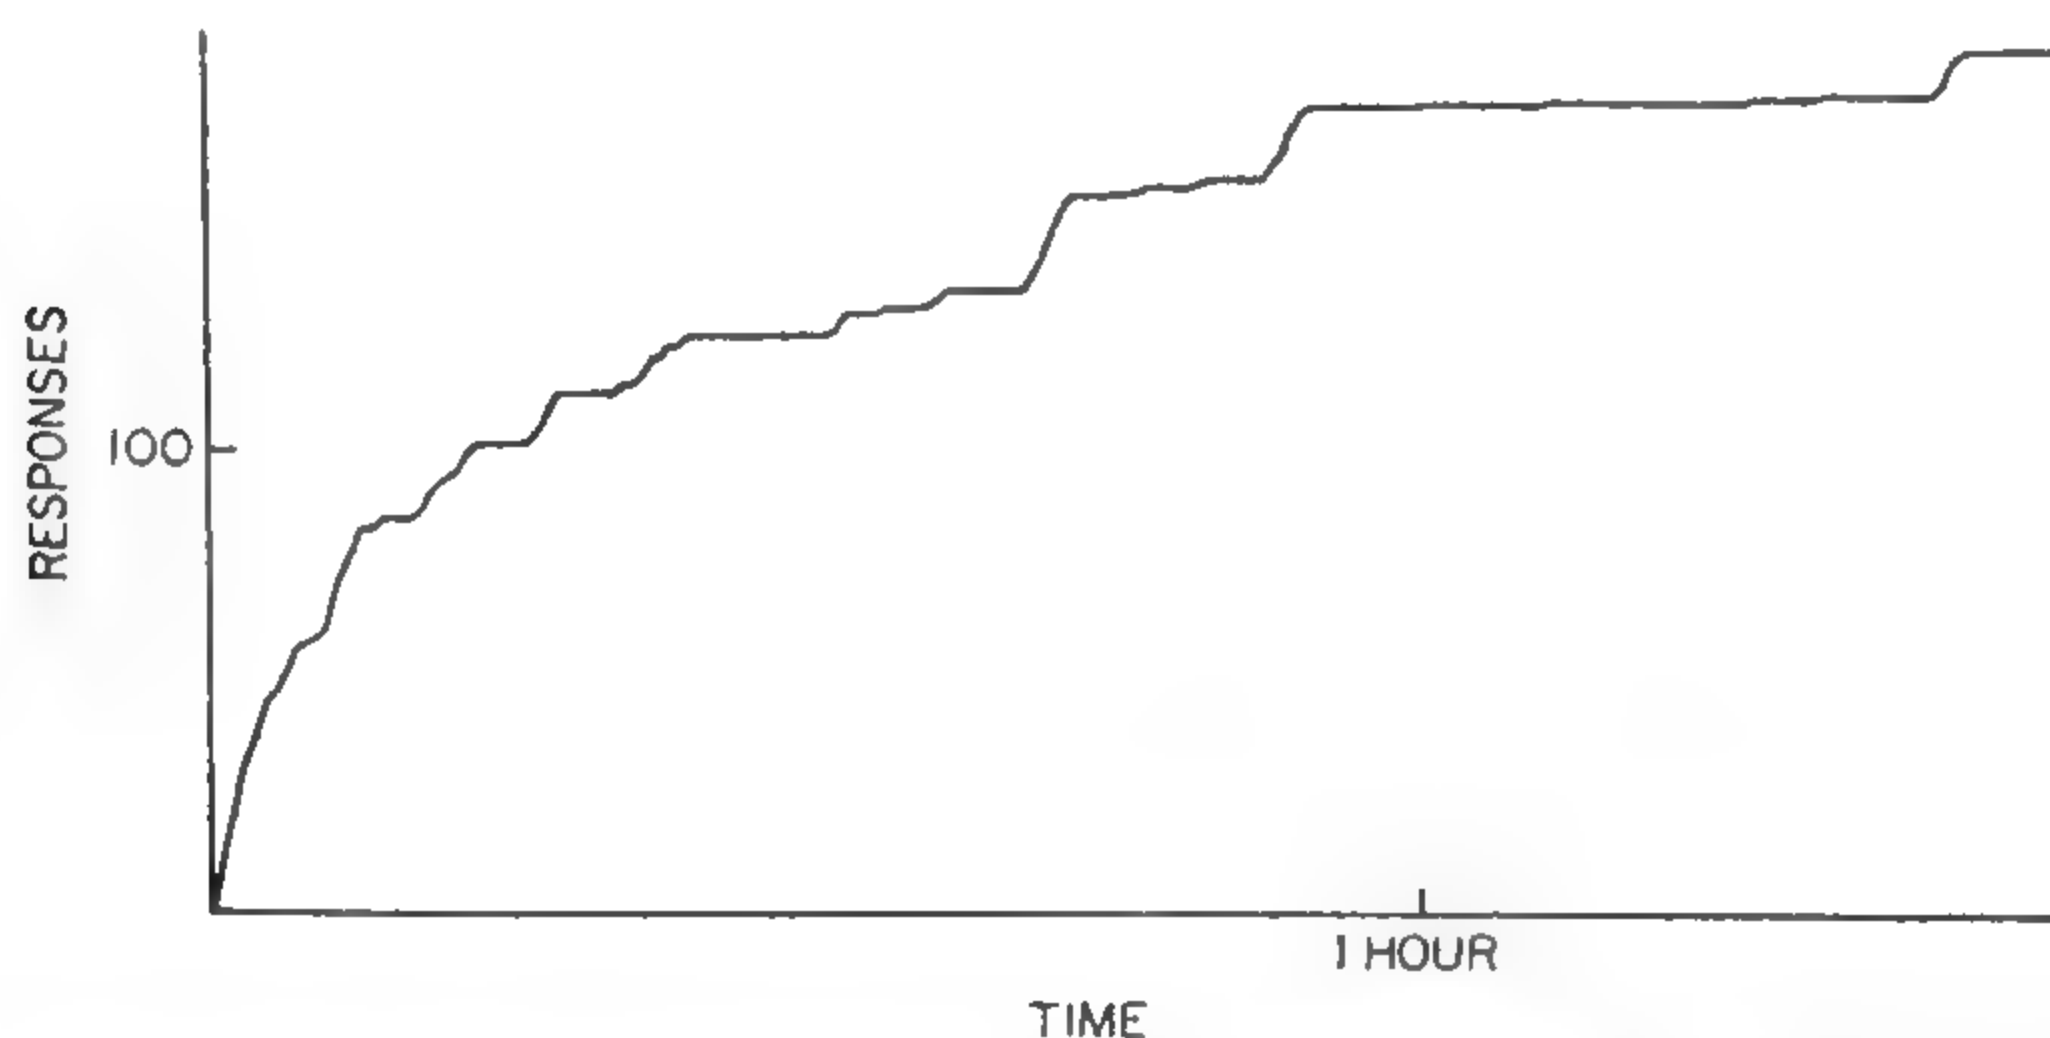


Figure 8. An extinction curve for the lever response in the Skinner box. The cumulative number of responses is plotted on the ordinate. The rate of responding is high at the beginning, but then it begins to decline and become more irregular. Long periods of time pass without responding, until, finally, the animal almost completely stops responding. (Skinner, 1938, data from F. S. Keller and A. Kerr.)

In the same way, the classical conditioned response gradually grows weaker if the unconditioned stimulus is removed. Thus, in Pavlov's experiment presented in the first chapter, if the tuning fork is sounded a number of times without being paired with the meat powder, the salivary response to the tuning fork gets smaller and finally disappears.

Resistance to extinction. It is apparent from the two examples we have just looked at that reinforcement creates a reserve of responses that can be emitted by the animal after the reinforcement is withdrawn. This reserve, or potential for responding in the absence of reinforcement, is usually called *resistance to extinction* or, sometimes, the *reflex reserve* (Skinner, 1938).

Resistance to extinction depends, to some extent, upon the number of reinforcements given an animal before extinction is begun. Figure 9 shows that if a small number of reinforcements is given, few bar-pressing responses are emitted in extinction before animals cease responding altogether. As

the number of reinforced responses is increased, however, the number of responses in extinction also increases. But the number of responses increases only up to a limit. Thus, for low numbers of reinforcements, resistance to extinction depends upon the number of reinforcements that have been given before extinction is begun. Maximal resistance to extinction, however, is obtained with relatively few reinforcements, and after a certain point one can go on giving reinforcements without appreciably increasing resistance to extinction (Skinner, 1938).

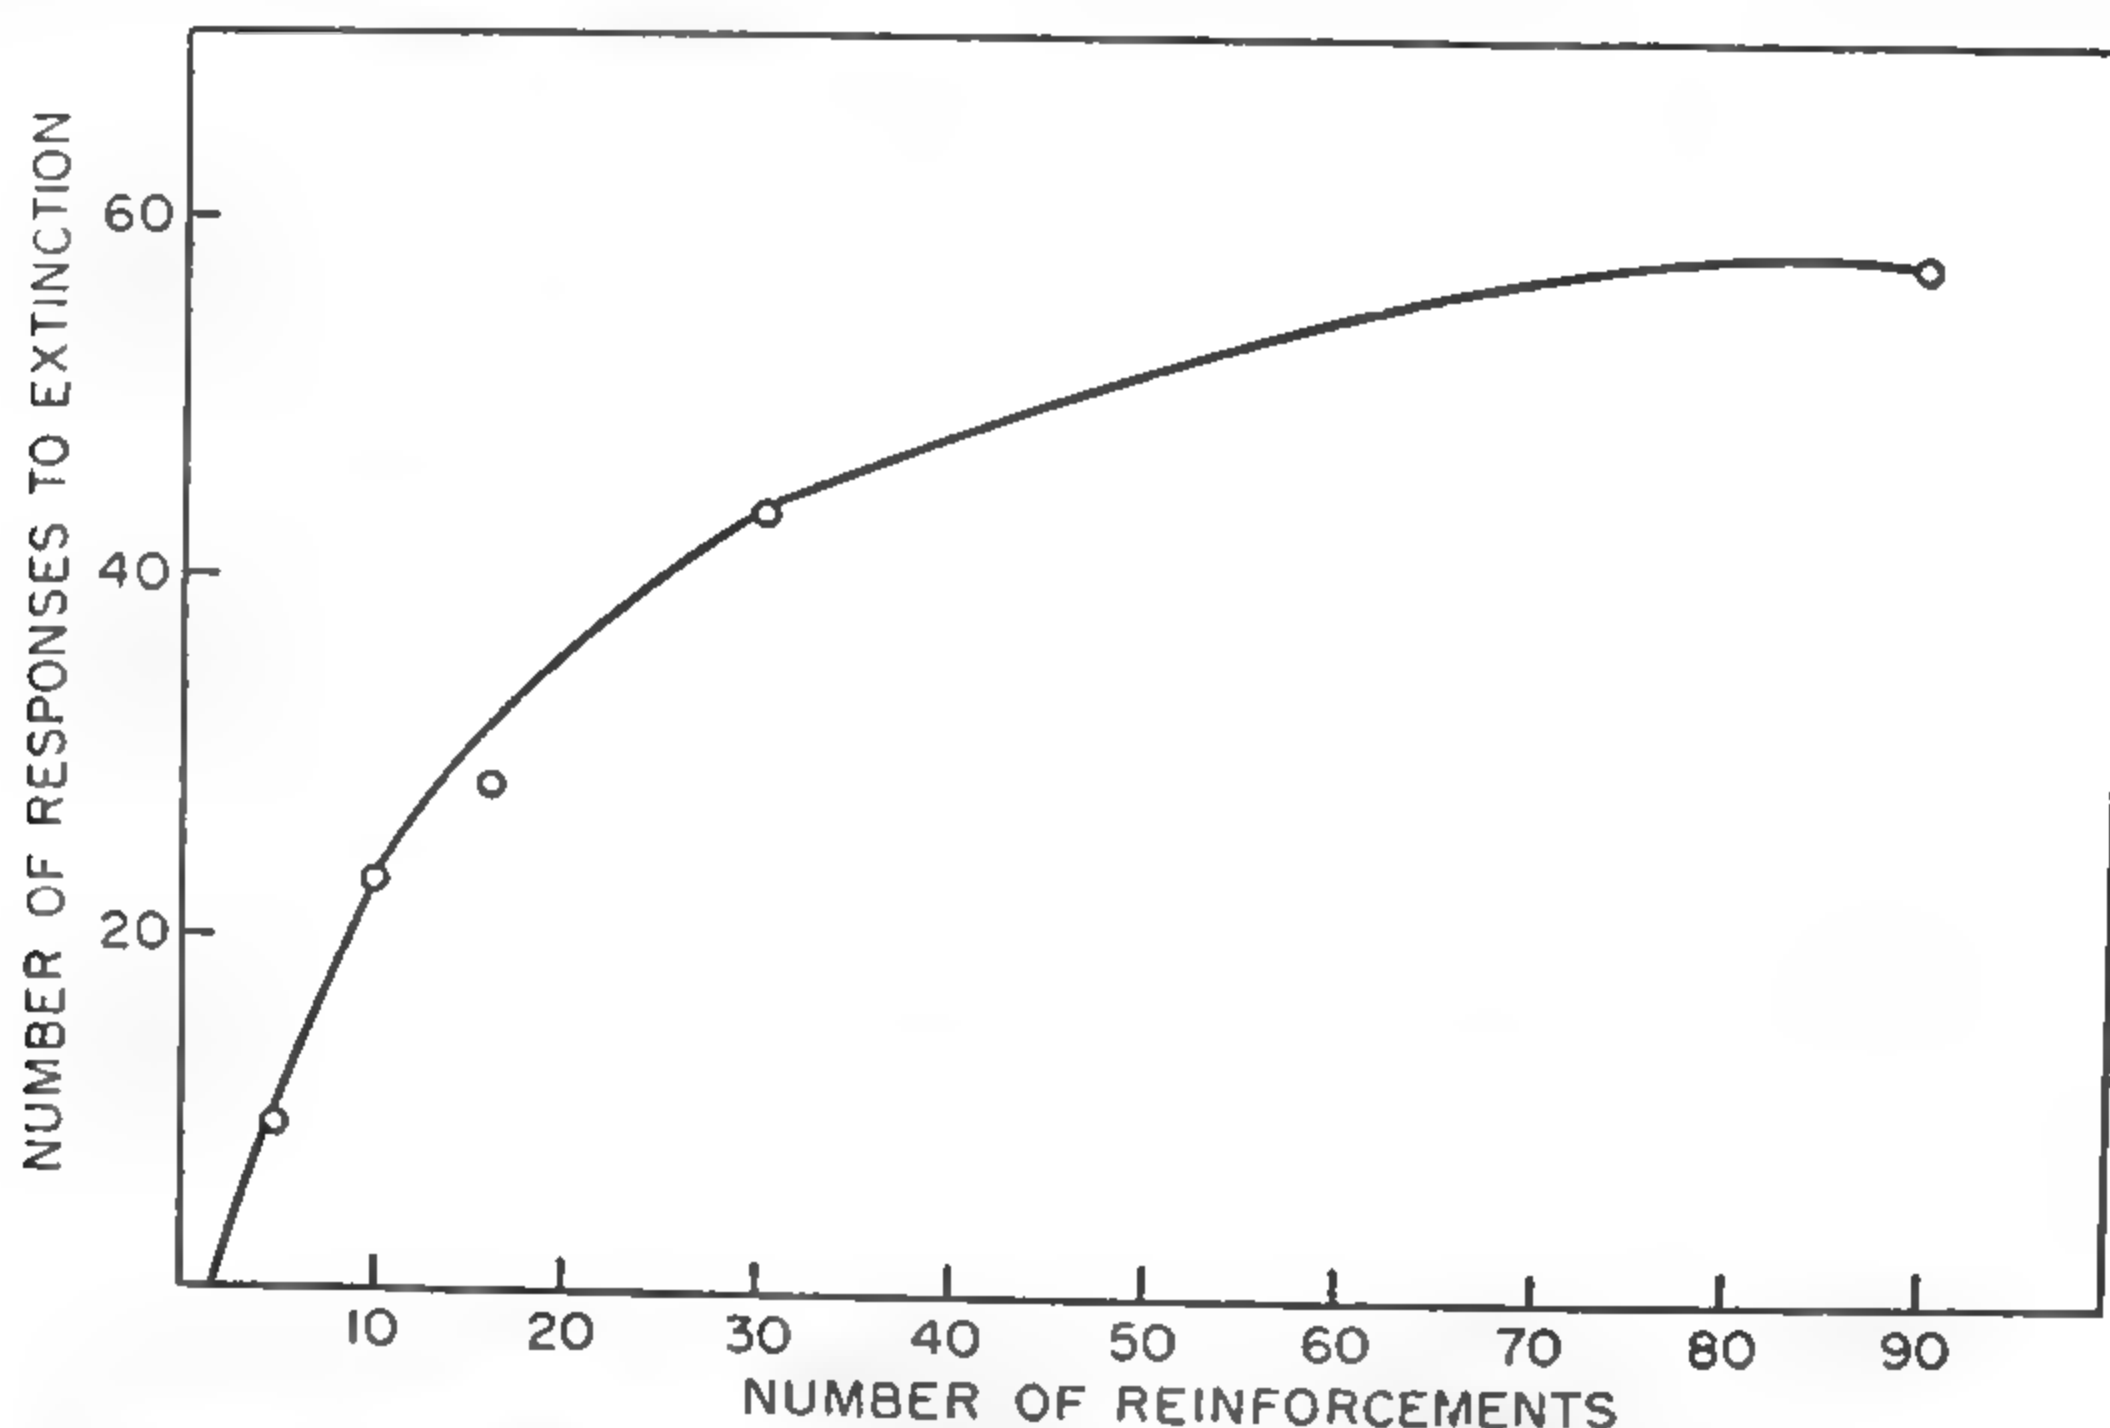


Figure 9. Number of lever responses during extinction as a function of the number of reinforcements administered during training. (Redrawn from Perin, 1942; data from Perin, 1942, and Williams, 1938.)

Time and Effort in Extinction

Two variables of considerable theoretical importance in extinction are (1) the distribution of responses in time and (2) the effort required to make responses during extinction. These are linked here because together they provide information important to the theory of extinction. Let us first look at the problem of the distribution of responses.

Distribution of responses in time. One of the main advantages of giving discrete trials during an experiment on learning is that the experimenter can control the rate at which animals respond. In the Skinner box, as it is usually used, the rate of performance depends upon the animal and such indirect means of control as hunger. Many experimental examples of instrumental learning, including bar pressing, can be so arranged that discrete trials are given after the fashion of the classical conditioned response. For ex-

ample, it is a simple matter to withdraw the bar from the Skinner box and to present it again only when we wish to begin the next trial.

Control over the rate at which responses occur is important to students of learning because Pavlov (1927) reported that a classical conditioned response extinguishes much more rapidly when trials occur close together than when they are spread apart. More recently experiments have been performed to see if this is also true of instrumental conditioning.

First of all, it seems clear that performance of instrumental behavior during *reinforcement* is impaired when trials are massed together (Hovland, 1936; Spence and Norris, 1950). Secondly, most studies show that massing trials during extinction has the effect of decreasing resistance to extinction or of decreasing the tendency to respond on any given trial during extinction (Gagne, 1941; Reynolds, 1945a; Rohrer, 1947; Teichner, 1952). Other studies, however, show that under some conditions resistance to extinction can be greater with massed trials (Sheffield, 1950) and that under still other conditions the spacing of trials probably makes no difference (Porter, 1939).

Teichner (1952) points out the probable reason for these differences. He suggests that, in general, acquisition is slower when trials are massed closely together and that extinction is faster under these conditions, but that these effects are relatively small and easily obscured by other conditions. For example, Teichner's experimental results on a running response in rats show that the effect of spacing trials during extinction depends upon their spacing during acquisition. The greatest resistance to extinction is found when the same spacing of trials is used for both; when different spacing is used, resistance to extinction decreases. Since Stanley (1952) shows experimental evidence for this same point, we can conclude that maximal resistance to extinction is produced by using the same distribution of trials in both extinction and conditioning and that slightly greater reduction in resistance to extinction occurs if the distribution of trials in extinction is faster, rather than slower, than that in conditioning.

Effort and extinction. The role that effort plays in extinction is important for theoretical reasons. One of the major theories of extinction, in fact, has assigned the decrement in response to a fatiguelike inhibition created by continued responding.

A well-known experiment by Mowrer and Jones (1943) provides evidence suggesting that effort does influence resistance to extinction. These investigators trained rats to press a lever for food in the Skinner box. During acquisition the lever was counterweighted in different amounts, so that the

rats learned to press the bar when it required 5 grams to activate the feeder, when it required 42.5 grams, and when it required 80 grams. At the end of acquisition the rats were divided into three groups, and each group was extinguished with a different counterweight attached to the lever. There was a simple negative relationship between resistance to extinction and the amount of effort required to push the bar. This general effect is confirmed in a study of Solomon's (1948) in which the distance rats had to jump to reach a goal box was the variable producing effort in the response.

These two experiments leave something to be desired, however. It is possible that the differences in resistance to extinction were due to variations in the amount of training given the different responses. When careful controls are introduced into an experiment on effort, resistance to extinction appears to vary only a small amount for responses requiring different amounts of work (Maatsch, Adelman, and Denny, 1954). Furthermore, the differences that do occur may be the result of differences in operant level of responding for various conditions; that is, rats may "spontaneously" engage in easy responses more frequently than in hard ones.

Thus, while it seems likely that work and time are variables in determining the course of extinction, they are not the most important ones. The results achieved in well-controlled experiments suggest that other factors are much more significant in determining the behavior of individual animals. As we shall see presently, this conclusion has important implications for theories of extinction.

Spontaneous recovery and disinhibition. Pavlov, in his studies of conditioned salivation, discovered two phenomena that he thought to be of considerable theoretical importance. They are called *spontaneous recovery* and *disinhibition*. Again, we link them together because they are important in the theory of extinction. In both spontaneous recovery and disinhibition the major effect is a recovery of the conditioned response after a certain amount of extinction. Thus, a strong argument can be made that extinction generates an active inhibition of the conditioned response and that this inhibition can be removed by particular conditions. Let us examine spontaneous recovery rather closely, since it is the more important and pervasive of the two phenomena.

Spontaneous recovery. An example from Pavlov's laboratory will illustrate the basic principle of spontaneous recovery. Pavlov (1927) produced a conditioned salivary response by allowing a dog first to *see* some meat powder and then to eat it. After the dog had learned to salivate at the sight of the meat powder, Pavlov extinguished the response by no longer putting

the meat powder into the animal's mouth. It took only a few trials of this sort to extinguish the conditioned response. The animal was then removed from the experimental room for approximately two hours. When the dog was brought back and allowed to look at the meat powder again, the conditioned response returned. True, the amount of saliva flow was not so great this time as it had been before the extinction process was begun, but there was no doubt that the conditioned response had recovered to a considerable extent—actually about one-sixth of the original amount.

Another experiment (Ellson, 1938) illustrates the occurrence of spontaneous recovery in an instrumental conditioned response. Four groups of rats were trained to press a lever in the Skinner box. Following the conditioning procedure, all animals were permitted to operate the lever without reinforcement until they had failed to respond for a period of five minutes. One group of rats was then removed from the box and allowed to rest for five minutes. These rats were replaced in the box and additional unreinforced responses were counted. The second group of rats resumed extinction after a pause of 25 minutes, the third group after 65 minutes, and the fourth after 185 minutes. All animals showed an immediate tendency to press the lever when replaced in the box, even though all of them had failed to press the lever for a period of five minutes prior to their removal from the box. Figure 10 shows that the amount of recovery, or the number of responses after rest, is an increasing function of the length of the rest period. Like most effects associated with the learning process, however, this function

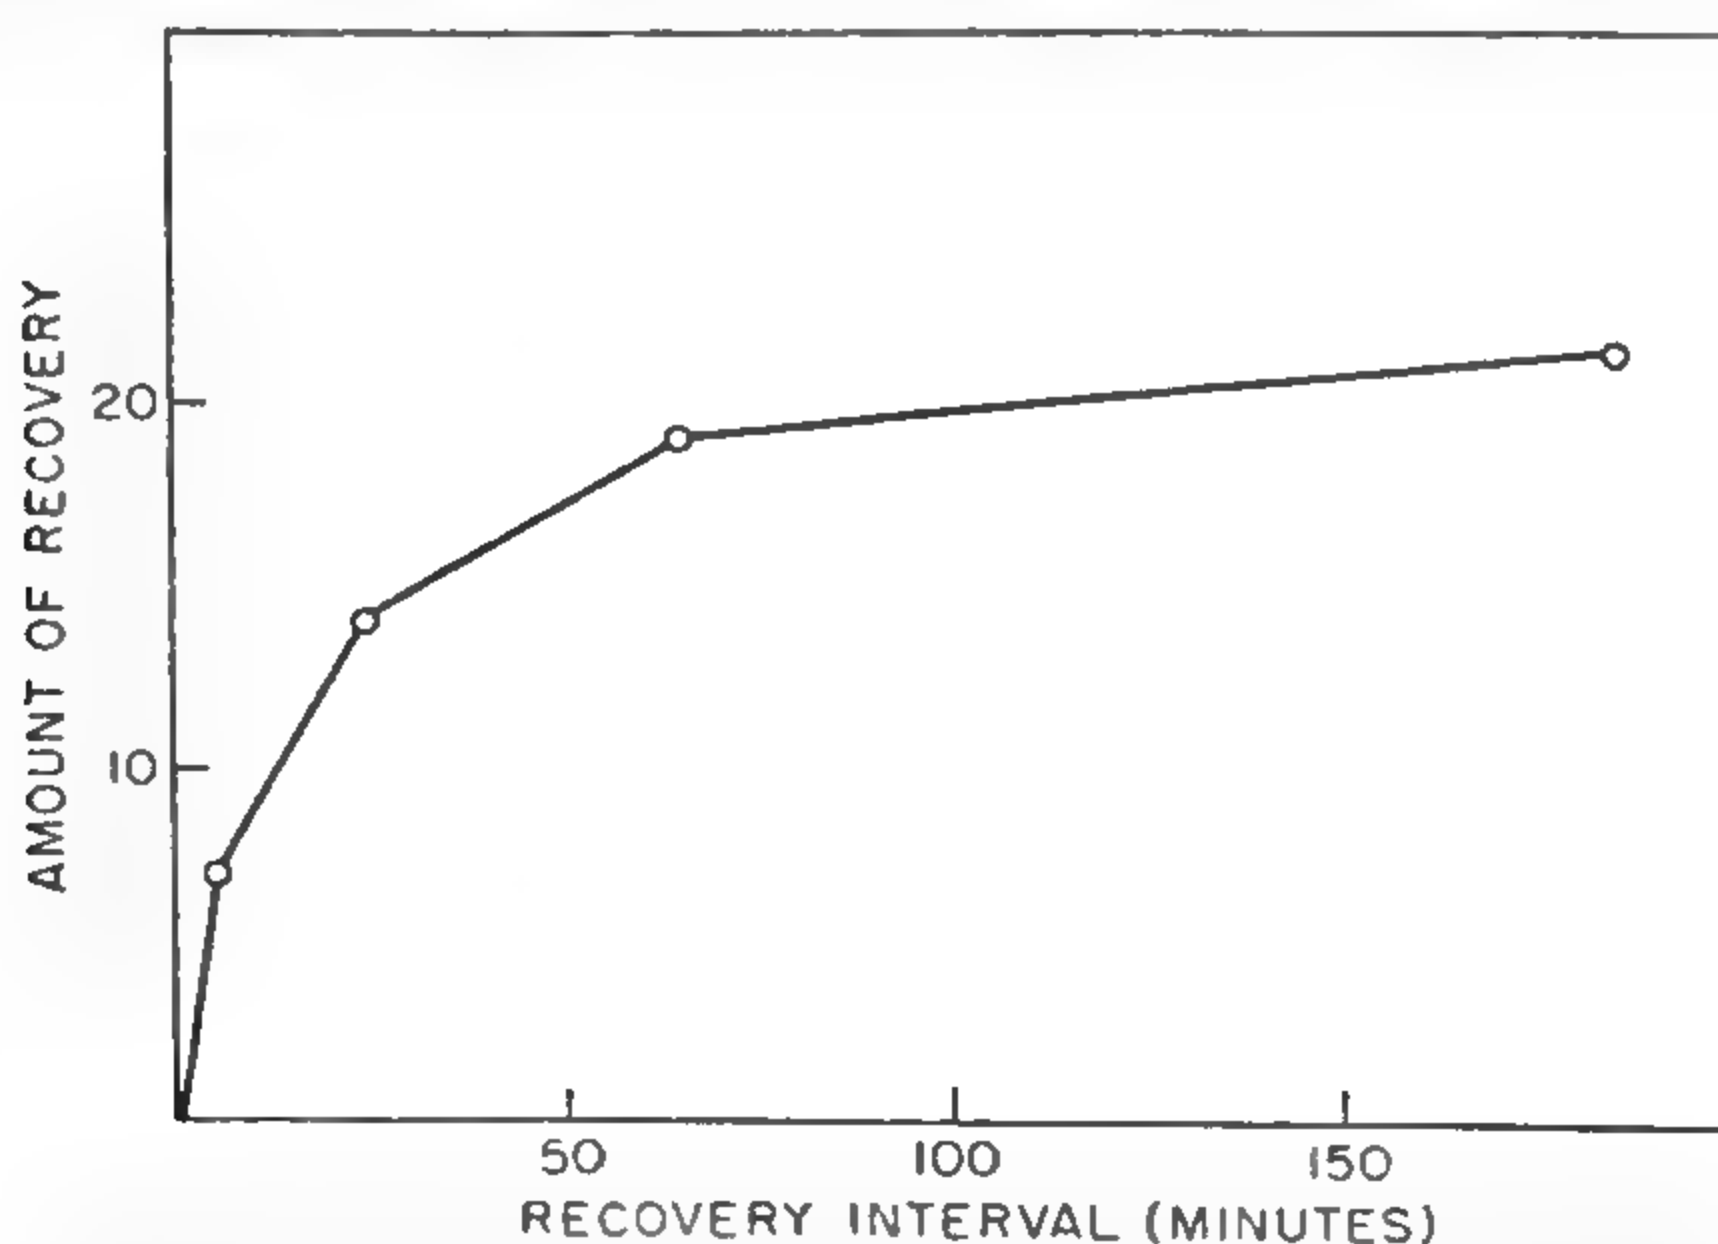


Figure 10. Number of lever responses in a second extinction period following various intervals of rest. In this study most spontaneous recovery occurs within an hour after original extinction. (Data from Ellson, 1938.)

seems to have its limit, so that an extremely long rest period yields no more spontaneous recovery than a moderately long one. In this experiment the amount of recovery reached its limit at about one-fifth the original strength of the response.

Spontaneous recovery is a genuine and stable part of the extinction process. We may, for example, try to stop our dog from begging at the dining-room table by ignoring it when it does so, and we may find by the end of the meal that we have been successful—the dog has retired to the living room. Nevertheless, the next evening it will probably beg as before. Since spontaneous recovery is incomplete, however, the dog will not be so persistent this time and will stop sooner. Repeated extinction periods will finally eliminate the response.

Disinhibition. The term disinhibition stems from Pavlov's theory of extinction, and he placed great stress on the importance of this effect. In his original demonstration of disinhibition, Pavlov established a conditioned response to the sight of meat powder. Once it was well established, he extinguished it by showing the meat to the dog without letting the animal eat it. When the conditioned response was completely extinguished, Pavlov paired a new stimulus with the conditioned stimulus. This stimulus was not one that could elicit the conditioned response in itself, but when it was paired with the *extinguished conditioned stimulus* it did. Thus, the salivary response to the sight of food, which had been completely extinguished, came back again when the new stimulus, vibration, was presented with the conditioned stimulus.

THEORIES OF EXTINCTION

Animals cease responding when they are no longer rewarded. So simple a fact would hardly seem to demand very subtle theoretical interpretation; yet extinction has been one of the aspects of behavior involving the most complicated theorizing. Actually, of course, a theory of extinction is fundamental to any theory of behavior, for the fact that animals cease to perform unreinforced acts is just as important as the fact that they learn reinforced acts in the first place.

In the following section we shall study a few of the more important theories of extinction. Then we shall try to evaluate the current status of the theoretical questions and examine some of their weaknesses.

Response-produced inhibition theory. The response-produced inhibition theory is now the classical theory of extinction. Like "classical" theories

in many other fields, however, it has been thoroughly demolished by critics and the rubble used to build new theories. Although the response-produced inhibition theory is no longer acceptable in its entirety, two of its ideas are still of considerable importance. For this reason, and also because it provides a good case history of the fate of psychological theories, we shall look at it in some detail.

The most complete statement of the response-produced inhibition theory comes from Hull (1943), though he took the basic concept of inhibition from Pavlov and borrowed other ideas from a variety of sources.

Reactive inhibition. Hull started with the assumption that whenever an organism makes a response it also generates some inhibition to that response. This means that, all other things being equal, once a response takes place, it is somewhat less likely to occur in the immediate future. Such inhibition may be thought of as analogous to fatigue in that it makes the next response more difficult. Like fatigue, this reactive inhibition disappears after a period of rest; if there is not enough rest between responses, however, the inhibition accumulates from response to response.

Hull (1943) advanced the hypothesis that the amount of inhibition that results from a series of responses is a positively accelerated function of the amount of work the response involves.¹ In other words, as the response requires more effort, the inhibition accumulates at an increasing rate. Further, he specified, rest reduces the inhibition generated by each response in a way that is best described by a simple negative exponential function. Eventually, with enough rest, the response will completely recover its strength.

In this theory of Hull's, reactive inhibition will accumulate whether the response is reinforced or not. If the response is reinforced, the positive effects of reinforcement must overcome the negative effects of reactive inhibition. Most of the time the positive effects will keep the animal responding, but there is an interesting exception—a matter of experimental fact—which lends plausibility to Hull's notion. If, during conditioning, trials are massed very close together, the animals may slow their rates of responding and even stop altogether, even though the reinforcement is still present (Hovland, 1936; Calvin, et al., 1956). This effect has been called "inhibition of reinforcement."

¹ Hull later (1951) considerably revised the substance of this assumption; he suggested that for responses involving large amounts of work (relative to the ability of the animal involved) the inhibition generated was not much greater than that for smaller amounts of work. In other words, the function which related responding and work was S-shaped.

The most important aspect of the response-induced inhibition theory, however, concerns the results of extinction. When responses are no longer reinforced, reactive inhibition is thought to accumulate without being counteracted by any positive effects of reinforcement. The result is that the animal eventually ceases to respond. Since reactive inhibition disappears with time, we should predict spontaneous recovery of the response after a rest. Remember, however, that in the example from Pavlov's laboratory, spontaneous recovery was only about one-sixth the original response strength. At best, spontaneous recovery regains approximately 50 per cent of the response strength at the end of conditioning. Since spontaneous recovery is incomplete, Hull postulated a second factor, *conditioned inhibition*, to account for the permanent effects of extinction.

Conditioned inhibition. Hull regarded reactive inhibition as a negative motivational state. Presumably, animals are motivated to avoid reactive inhibition as much as they are motivated to avoid an electric shock. If we accept the need-reduction theory of reinforcement, it is easy to see how escape from, or avoidance of, reactive inhibition would constitute a reinforcement. Thus, when an animal ceases to respond, its fatigue state, or reactive inhibition, is reduced. This, in keeping with the need-reduction notion, is reinforcing. Reduction of reactive inhibition supposedly reinforces the animal for doing nothing, and the resulting failure to respond is called conditioned inhibition.

The notion of conditioned inhibition does not entirely depend upon the need-reduction theory of reinforcement, even though Hull presents the idea in this context. Reactive inhibition could bring about failure of response. Simply by contiguity, failing to respond could become conditioned to the cues that previously brought about the response. Thus, learning not to respond could as easily be the result of simple association as of reinforcement by need reduction.

How reactive and conditioned inhibition work in extinction. The theory of extinction based on reactive and conditioned inhibition seems complicated, but its essentials are simple. Extinction is the result of an active inhibition of a learned response. This inhibition is composed of two parts—reactive inhibition, which recovers with rest, and conditioned inhibition, which does not. Spontaneous recovery of an extinguished conditioned response is predicted from the assumption that reactive inhibition disappears with rest, while the permanent part of extinction is assumed to be the result of conditioned inhibition. The theory predicts more rapid extinction when

trials are massed during extinction and when a greater amount of effort is required in performing the conditioned response.

As we shall see, this two-factor theory of extinction (as it has been called) has been applied to a variety of problems in behavior—for example, to the fact that rats generally prefer not to repeat the same response on successive trials in a maze. To take another important instance, it has been applied to the pronounced effects of massed practice and recovery with rest on the performance of complicated human skills, such as learning to type. Despite the widespread application of the theory, however, we now suspect that it contains some serious defects, and it is therefore important that we look at some of the criticisms of the theory. Since attempts to examine this theory experimentally have led to some valuable new information about behavior, the next section will serve as more than just a critical analysis of the response-produced inhibition theory.

A critical analysis of the response-produced inhibition theory. The response-produced inhibition theory, as stated by Hull, has been criticized on three grounds: (1) the theory is a good deal more ambiguous than it first appears (Koch, 1954); (2) it leads to contradictory deductions (Gleitman, Nachmias, and Neisser, 1954); and (3) the empirical evidence for it is weak, and that against it, strong.

Let us first look at the empirical evidence against the theory. This is perhaps the least critical objection to the theory, since almost any theory can be modified to fit a new fact. The facts stated below, however, are particularly awkward for the response-induced inhibition theory.

Extinction without responding. The most important single idea in the response-produced inhibition theory is that decrement in responding is the result of a negative feedback from responding itself. Thus, the more effort a given response requires, the greater the difficulty in producing an additional response. Yet there is evidence that a decrement in responding can occur even when no overt response causes it. There is the phenomenon of “subzero” extinction, reported by Pavlov (1927). Pavlov showed that if we extinguish a classical conditioned response to the point at which it completely disappears, and then *continue* to present the conditioned stimulus, these additional trials weaken the conditioned response even further (shown by spontaneous-recovery tests). This means that presenting the conditioned stimulus alone was effective in producing further reduction of response strength—it was not necessary for the animal to give the conditioned response. Thus, response-produced inhibition could not have occurred.

It is also difficult for the response-induced inhibition theory to explain data derived from experiments on nonresponse extinction. The first example of such extinction comes from Seward and Levy (1949). They trained rats to run down a runway to secure food placed in a goal box at the other end. After conditioning, one group of rats was placed directly in the goal box when it was empty of food and permitted to see that the reinforcement was no longer there. A second group of rats did not have the benefit of this preview. The next day, all the rats were given a series of extinction trials. The group that had been able to see beforehand that the food was no longer there extinguished approximately twice as quickly as the other rats. It was evident that the response had been reduced in strength by the sight of the empty box and therefore extinction took place without the running response.

Several other experimenters have been able to find "extinction without responding" (Deese, 1951; Hurwitz, 1955; Moltz, 1955), and others have not (Bugelski, Coyer, and Rogers, 1952; Sharlock, 1954); the effect may depend on factors that are not yet well understood. Those experiments that show nonresponse extinction add some interesting information to our knowledge of the extinction process.

In one experiment (Deese, 1951) rats were trained to go to one side of a single-choice T maze. After the animals had been trained, half of them were placed directly in the empty goal box to let them see that there was no food. All the animals were then given a series of extinction trials. As in the previous experiment, the animals that had been allowed to see that food was missing extinguished more rapidly.

In addition, a test was made for spontaneous recovery from nonresponse extinction. One-half of the animals permitted to see the empty goal box were given their extinction trials immediately after the look into the goal box; the other half was tested 24 hours later. If there were spontaneous recovery from the inspection of the empty goal box, the group tested 24 hours later should have been more inclined to choose the goal-box side of the T. However, there were no differences between these groups. A test of spontaneous recovery from ordinary extinction in the same maze showed that there was indeed spontaneous recovery when the rats had to run through the maze to get to the empty goal box (Figure 11). These results suggest that there was little or no response-induced inhibition developed by the extinction without responding, while the spontaneous recovery from extinction *by responding* might reflect the presence of inhibition induced by the running response.

The experiment by Moltz (1955) suggests that extinction without responding is the result of the reduction of the secondary reinforcement value of the cues from the reinforced side of the maze. This leads to the proposition that extinction involves the reduction in the strength of secondary cues to "command" behavior when they are no longer associated with primary reinforcements. Hurwitz (1955) points out that, in the Skinner box, if one extinguishes early members of the response chain (approaching the lever, for example), one also extinguishes the later members (pressing the lever),

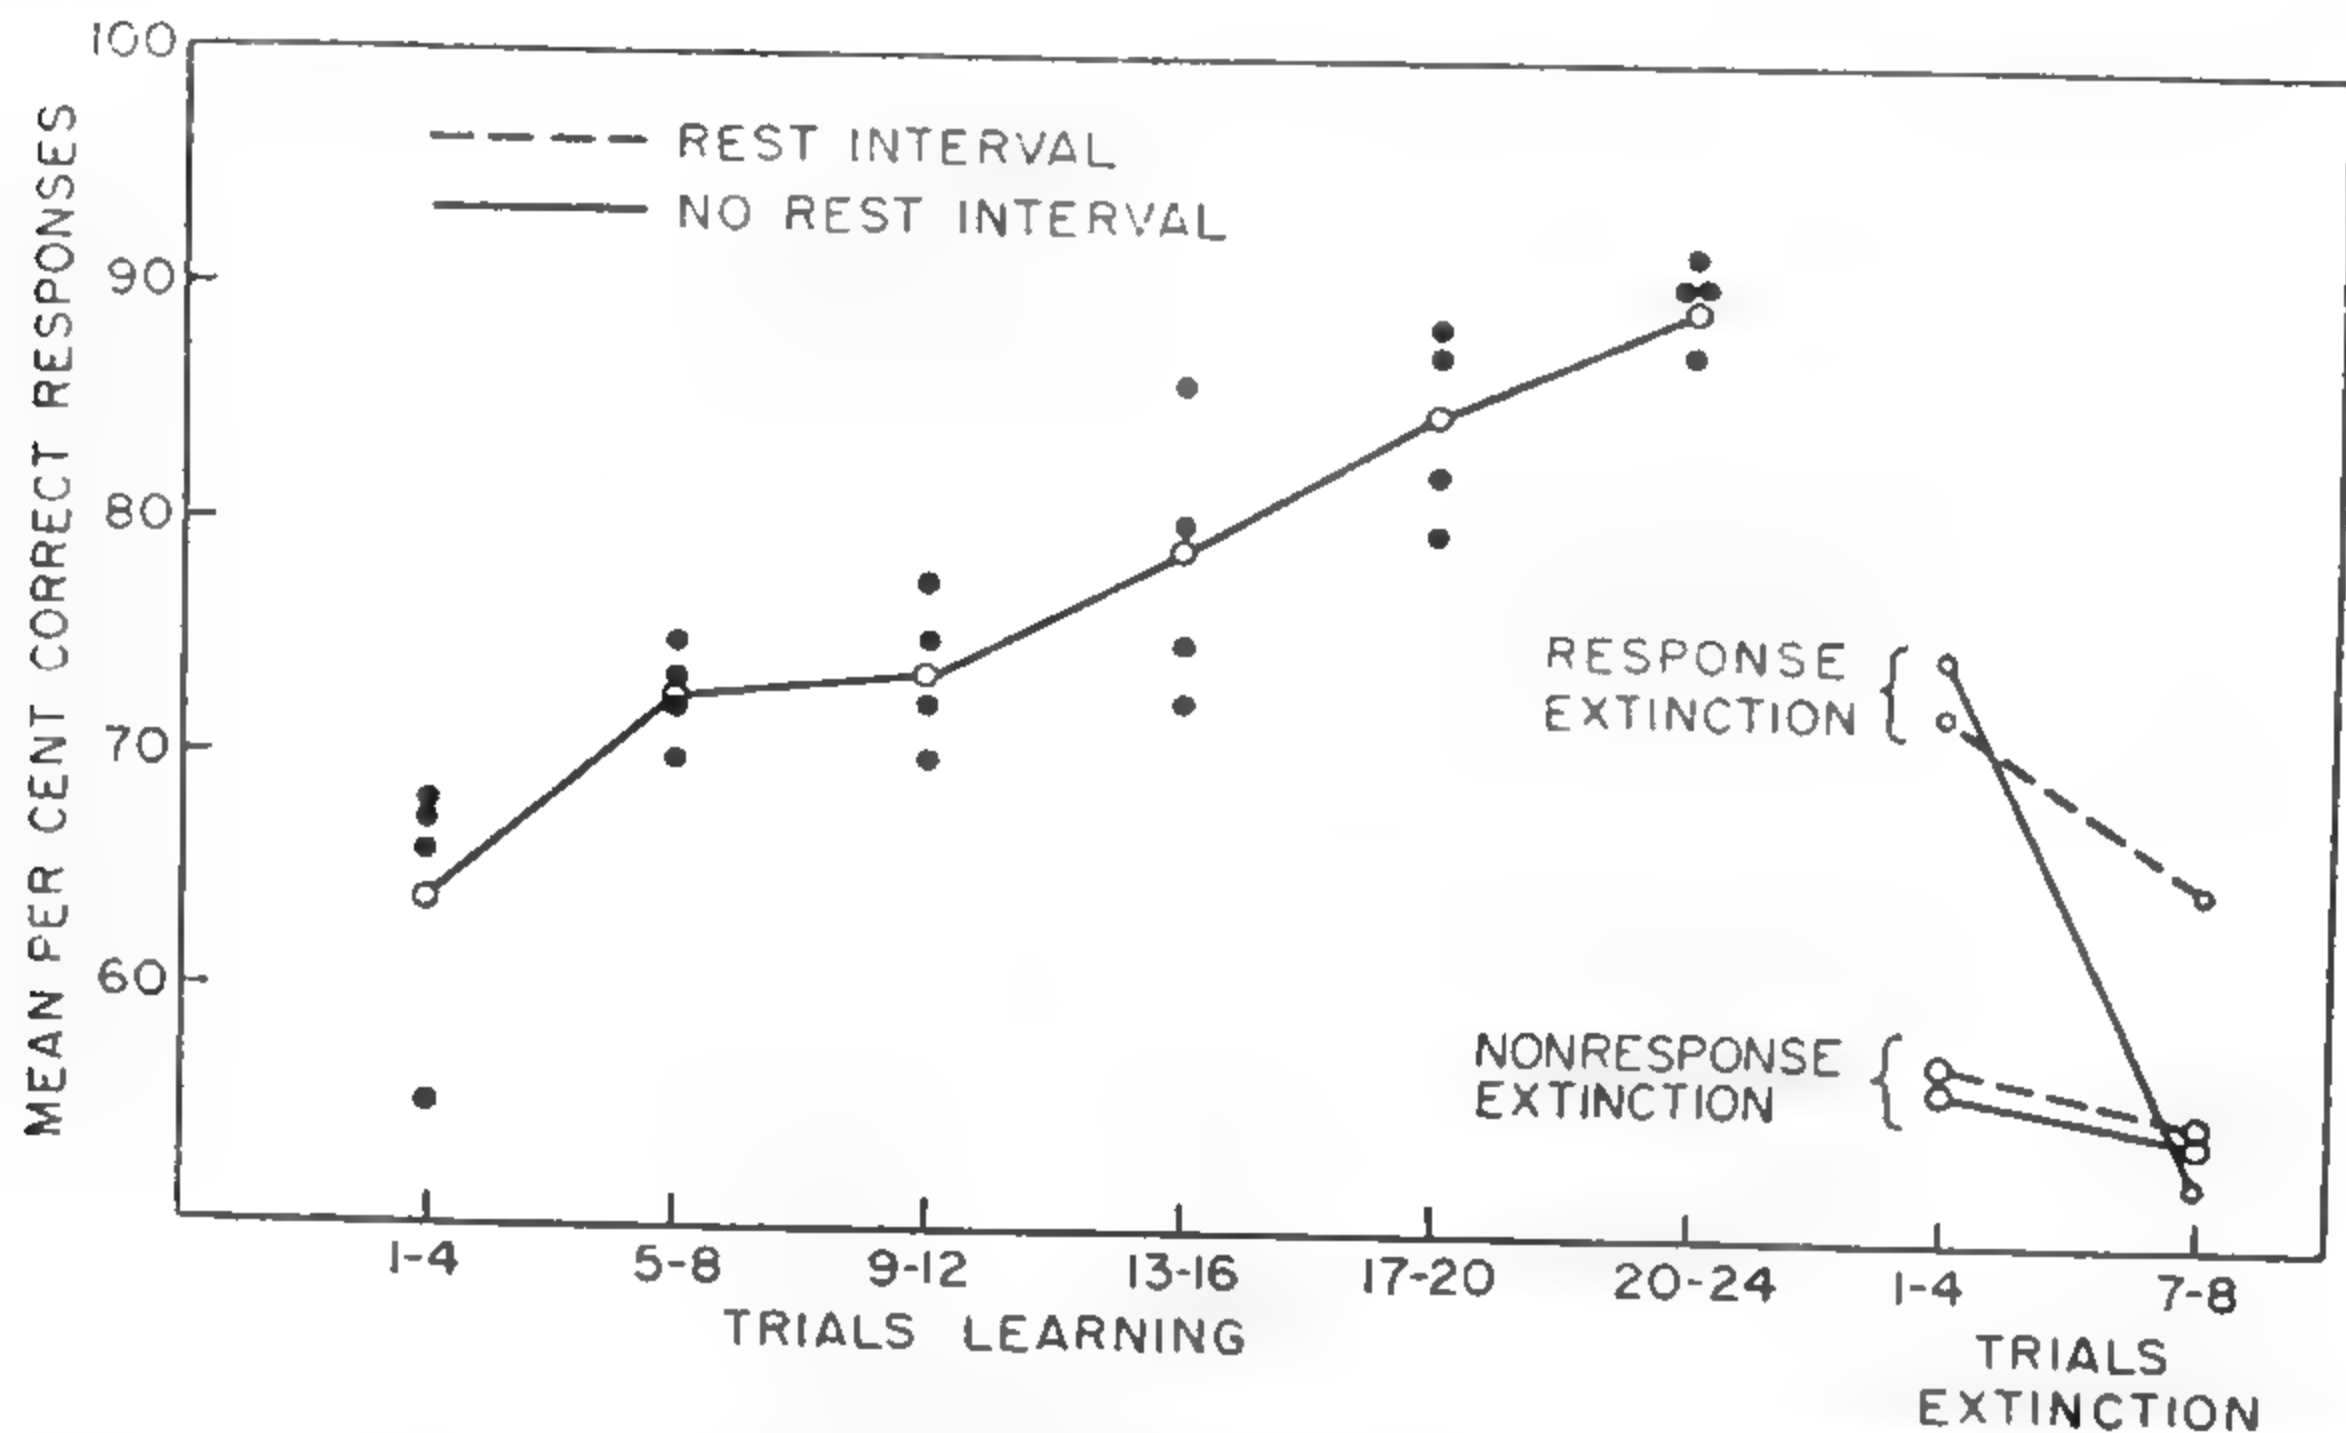


Figure 11. Learning and extinction of the correct choice in a T maze. "Response" extinction indicates tests for the number of correct choices made after the usual extinction procedure. "Nonresponse" extinction indicates choices made after the animals had been placed directly in the empty goal box. Running to the goal box does not seem necessary to produce extinction. There is spontaneous recovery following "response" extinction but little or none following "nonresponse" extinction. (Deese, 1951.)

and he suggests that this may be due to the weakening of the goal gradient or gradient of reinforcement.

Response inhibition and stimulus satiation. One of the facts that have tended to support the theory of response-produced inhibition is provided by studies which show that, given a free choice, animals prefer not to repeat responses they have just performed. For example, if rats are not reinforced for going to either side of a single-choice T maze, they will tend to alternate choices and not repeat their last choice. If response-produced inhibition provides an explanation of alternation behavior, there are a number of deductions that can be made from the properties of response-produced inhibi-

tion. For example, decreasing the time interval between trials ought to increase the frequency of alternation. Likewise, increasing the effort involved in making a response should increase alternation. In another instance, forcing animals to go consistently to one side of a T maze for a while ought to make them tend to choose the opposite side. This tendency to choose the opposite should decline if an interval of time is allowed between the forced trials and the free-choice trials (reactive inhibition disappears with rest). In general, these deductions stand up to experimental test (Heathers, 1940; Zeaman and House, 1951; Riley and Shapiro, 1952).

Despite this rather impressive array of evidence, investigators of animal behavior, other than learning theorists, have long held another explanation for the failure of responses with repeated elicitation. One investigator (Pracht, 1953) studying the gaping response (the yawnlike opening of the mouth observed in young birds when the parent bird returns to the nest with food) showed that, while repeated elicitations of the response will exhaust it, changing the stimulus will make the response come back again. Actually, many experiments have shown the major cause of alternation and related phenomena to lie in stimulus conditions rather than in the repetition of responses per se.

One simple experimental result seems strongly against a theory of alternation exclusively in terms of response-produced inhibition. If we exchange the stimuli associated with each pathway in a two-alternative maze, rats will tend to repeat rather than alternate their choice of pathways (Glanzer, 1953). This suggests that the rats go where there are "new" stimuli. Similarly, another experiment shows that rats alternate places they go to more frequently than specific responses in getting there (Montgomery, 1952).

A rather exhaustive experiment by Rothkopf and Zeaman (1952) shows that the control of alternation in a T maze is not simple; they suggest that both inhibition from response and such factors as novelty of external stimulation control the tendency to alternate. Thus, while it is probable that response-induced inhibition plays a part in such phenomena as alternation of responses, it is certainly not the only factor, and perhaps not even the most important.

The present status of the response-produced inhibition theory. Putting all the evidence together—evidence concerning the distribution of trials, the effects of effort on response tendency, alternation behavior, etc.—we must conclude that responding has an inhibitory effect upon future tendency to respond. Beyond this, however, we cannot go, for it is equally evident that this feedback from responding is not necessary to extinction. Furthermore,

it appears that there are rather severe logical problems in applying the notion of response-induced inhibition to extinction itself.

The most serious logical problem in the response-induced inhibition theory as it is stated by Hull is that it implies that *all* behavior must eventually cease, whether it is reinforced or not (Gleitman, Nachmias, and Neisser, 1954). For even though habit strength for any particular response will eventually reach a maximal value, inhibition continues to be added for each response. As inhibition is added, so is conditioned inhibition, which thus continues to grow while habit strength remains constant. The net effect of this will be to make the response weaken and eventually disappear. It is true, as we saw above, that responses do sometimes weaken in strength even with continued reinforcement and, indeed, they may even eventually disappear (Calvin, et al., 1956), but this is the rare exception.

The root of the matter is that the concept of conditioned inhibition is weak. As originally stated in Hull's theory of learning, it leads to contradictory properties. We have seen that extinction can occur when conditioned inhibition based upon response-produced inhibition should be very weak. Therefore, it would seem that the notion of conditioned inhibition is superfluous. This would leave us with the idea that there is a temporary, fatigue-like decrement produced by continued responding. This inhibition may have something to do with the organism *temporarily* ceasing to respond, but it cannot have much to do with the *permanent* effects of extinction. It is clear that, even if we do accept the general proposition of response-produced inhibition, we will need some concept that will help explain the permanent effects of extinction.

CURRENT ALTERNATIVE THEORIES OF EXTINCTION

Competition theory. The best-known alternative theory of extinction is that put forward by Guthrie (1935, 1952) and, to some extent, by Estes (1950). In a word, this theory attributes extinction of a particular response to the acquisition of competing response tendencies. Thus, when an animal ceases to respond in a particular situation, it is because it has learned not to do so by an association between "not doing so" and the stimuli. In some cases the animal will learn alternative responses and in others it will learn to do nothing.

This leads Guthrie to say that there are three ways in which responses may be eliminated. The first method is to introduce the conditioned stimulus subliminally, so that it will not elicit a response, and then to increase its

strength gradually. This is what we do when we train a horse to saddle by first putting on only a light blanket and then gradually working up to full gear. The blanket is not a sufficient stimulus to set off bucking, and the horse is not disturbed by the gradual addition of heavier loads. An experiment by Kimble and Kendall (1953) shows that this method of "toleration" works quite well in a simple avoidance habit in rats. A second method is to repeat the conditioned stimulus until the original response is exhausted and the organism is too fatigued to respond. This is the bronco-busting technique and, incidentally, the special case handled quite well by the conditioned-inhibition notion. The third method is to present the conditioned stimulus when it is mechanically impossible for the organism to respond.

Aside from providing this common-sense analysis of methods of eliminating behavior, the theory of competition advanced by Estes has the merit of considerable logical elegance. For one thing, conditioning and extinction become simply different aspects of the same thing—they are both learning by a process of associating elements of particular responses with elements of stimulus complexes. Furthermore, as Estes (1955) points out, spontaneous recovery under this theory is not a special phenomenon of inhibition, but results from the fact that samples of stimuli conditioned (or extinguished) on any one trial may not be present on the next; thus the tendency to respond can "spontaneously" change between trials.²

There are difficulties, unfortunately, with the competition theory, too. For one thing it becomes rather cumbersome when applied to the general cases of indirect extinction discussed under Extinction without Responding. Furthermore, there are practically no simple relationships between the properties of conditioning and extinction, for a maximal rate of responding can sometimes be achieved in the Skinner box with only one reinforcement, while extinction after one reinforcement may require as many as 150 responses before the rate of responding is returned to the unconditioned level (Skinner, 1938). In classical conditioning, extinction is usually more rapid than the original conditioning.

These, however, are relatively minor objections, not nearly so serious as those to the theory making use of the concept of response-induced inhibition. There has been a serious effort, moreover, to examine some of the experimental implications of the theory (see, for example, Lauer and Estes, 1955). Perhaps the most serious objection to the theory is not a logical one. That is, there is some reason for believing that extinction, in

² For an experimental analysis of the role of spontaneous recovery in a formal competition theory, see Homme, 1956.

part at least, reflects factors other than loss of habit. Let us examine these in the context of other theories.

The frustration theory of extinction. Many experimenters have observed that when an instrumental habit is extinguished animals tend to become emotionally excited, as if they were frustrated by the withdrawal of reinforcement. Skinner (1938), for example, has noted this in rats undergoing extinction and has attributed to it the comparative irregularity of extinction curves. More specifically, some investigators have assumed that extinction gives rise to a frustration drive and that this accounts for some of the effects one finds in a study of extinction.

One simple runway experiment (Adelman and Maatsch, 1955) shows that if rats are allowed to escape from the empty goal box during extinction, they do not show any loss for the originally reinforced habit of running down the runway. Of course, if these animals are confined in the goal box, they do show a loss of response strength. One interpretation of this is that escaping from the frustrating conditions of extinction does not produce the same emotional aversion to the empty goal box as being confined in it.

Another series of experiments (Lambert and Solomon, 1952; Lambert, Lambert, and Watson, 1953) show an effect that is difficult to interpret except by assuming that extinction gives rise to a state of frustration. These experiments show that the closer an organism gets to the reinforcement before it is frustrated, the harder it is to extinguish a habit. In one of the experiments this effect was discovered by blocking rats at various points in a simple runway; the farther the block was from the goal box, the more quickly the rats gave up running. In the other experiment, essentially the same effect was found with a token-reward habit in young children.

The interpretation of these experiments is that frustration is greater when the organism is blocked just before reaching the goal box than when it is blocked farther away. This frustration heightens the animal's drive and, hence, its activity, so that it interferes with either the acquisition of a new habit or "learning not to respond."

It is probably true that the failure of reinforcement for many instrumental acts does produce frustration, and that this has special effects on behavior, as the above experiments show (also, see Stanley and Rowe, 1954). Frustration does not provide an explanation of extinction itself, however; it only suggests how the process of extinction can be modified by emotional activity.

Extinction as loss of incentive value. Strangely enough, the hypothesis that is probably the simplest has received the least attention. This is the no-

tion that extinction has little or nothing to do with habit or learning, but, rather, produces a reduction in response strength by changing the value of the incentive (Johdai, 1955). The effects of extinction, according to this argument, can be considered simply as the limiting case of the reduction in size of reinforcement after learning has taken place (Dufort and Kimble, 1956). Given the fact that behavior varies according to the size of reinforcement, it seems a theoretical extravagance, as Dufort and Kimble point out, to assume that completely different processes take place in extinction—reduction of reinforcement to zero after previous reinforcement.

In the last chapter we discussed another special case of the absence of reinforcement under latent learning. In the experiments on latent learning we saw that animals could learn something about a maze even though they were not specifically reinforced in the goal box. The fact that the animals had learned the maze, however, was not evident until some incentive was given them. In the same way, if we extinguish a maze habit by withdrawing reinforcement after the maze has been learned, the fact that animals will now make more errors (see Tolman and Honzik, 1930) does not mean that they have unlearned the correct path to the goal box. Rather, the rats appear to have no interest in getting to the goal box. There are, of course, many implications in considering extinction to be a special case of zero amount of reinforcement. Many of these implications were discussed in the previous chapter under amount of reinforcement.

The present state of theories about extinction. We have seen that there is a temporary negative feedback effect on responding resulting from previous responses. This effect provides the basis for the concept of response-produced, or reactive, inhibition. Reactive inhibition disappears with rest, however, and except for special cases, it is doubtful whether it has much to do with the more permanent effects of extinction. In addition there is an analogous phenomenon in stimulation (Hinde, 1954; Glanzer, 1951); that is, all other things being equal, animals tend to avoid stimuli they have just been subjected to. It is likely that this “stimulus satiation,” like reactive inhibition, is largely temporary in nature, though this issue has not been thoroughly explored.

Furthermore, responses sometimes disappear because organisms learn new responses that displace the old ones. In classical conditioning this may mean that the organism acquires a new response to stimuli already conditioned to another, incompatible response. In instrumental conditioning it may mean that the organism learns a new method for receiving the same

reinforcement or that it learns that a different set of stimuli is not correlated with the occurrence of reinforcement.

There is one important question about this matter of extinguishing an old response by learning competing activity. It is not clearly established that learning the new response necessarily results in "unlearning" the old one. The theories advanced by Guthrie and Estes suggest that this is so, but there are good grounds for assuming that it is not always true. It is clear, at least, that learning a new incompatible response does not always reduce the old response to its original, unlearned strength. Those theories that reduce all conditioning to the paradigm of classical conditioning would suggest that this is so simply because all of the cues that were conditioned to the old response are not conditioned to the new one, so that stimuli conditioned to the old response are left over.

At any rate, one of the things that should soon be settled is whether competing responses always result in the unlearning of the originally established response or simply have a temporary effect on performance. If the latter is true, when the competing response disappears, the original response should spontaneously recover in strength. There is some experimental evidence that the latter case is true at least in some situations (Whiting and Mowrer, 1943). This is a complicated question, and we shall see it again when we consider discrimination and conflict in animal behavior, and transfer in human verbal behavior.

In addition to the factor of competition, it is possible that a large number of the examples of extinction studied in the laboratory (as well as many problems in animal behavior in nature) are the result simply of the change in incentive or reinforcement. After our discussion in the last chapter, we should suspect that such examples of extinction primarily involve changes in what animals are willing to do rather than in what they learn. Thus, many reductions in strength of response by nonreinforcement involve reduction in the momentary tendency to respond rather than any basic habit patterns.

PARTIAL REINFORCEMENT

We have seen that learned behavior is maintained in high strength by reinforcement, and that if it is not reinforced, it declines in strength. This effect of reinforcement determines *what* animals learn, since they cannot learn in situations in which they do not have some experience; the occurrence of reinforcement determines the kinds of experiences organisms are

going to have. Thus far we have considered only two particular cases of reinforcement, one in which occurrence of a response is reinforced, and the other in which there is no reinforcement at all. It is obvious, however, that there may be intermediate cases. We could, for example, reinforce every other occurrence of a response. In general, any case in which we reinforce behavior some, but not all, of the time is an example of partial reinforcement. The schedule of reinforcement is an extremely important variable in the control of behavior, and it will be worthwhile to examine this problem in some detail.

SCHEDULES OF REINFORCEMENT

In general, there are two ways we can control the patterning of reinforcements and responses. In the first case we can vary the *time* between reinforcements irrespective of the number of responses between reinforcements. Thus, we may arbitrarily decide to reinforce once every five minutes, regardless of whether the animal responds one hundred times or only once during the five minutes. The other way in which we can change the pattern of reinforcements is to vary the *frequency* of responses between reinforcements. In this case, we can decide to reinforce every tenth response, whether 10 responses take five minutes or one. In using these two methods, we can alter the reinforcement schedule so that it is either haphazard or highly systematic. In other words, we can reinforce randomly after one, three, or five minutes, or reinforce systematically every minute or every three minutes. Since these different ways of scheduling reinforcements have somewhat different effects, let us now look at them briefly.

Fixed-interval reinforcement. Fixed-interval, or periodic, reinforcement occurs on a fixed time schedule. This kind of reinforcement has a particular and systematic effect on previously learned behavior. When an animal becomes accustomed to a fixed-interval reinforcement schedule for, say, pressing a lever, its behavior becomes stable. In general, the animal will emit a constant number of responses per reinforcement. If, for example, the animal gives an average of 20 responses for each reinforcement and it is being reinforced once a minute, its rate of responding will be 20 responses per minute. If, however, the animal is reinforced once every two minutes, these 20 responses will be spread out over that period, so that its rate of responding will be only 10 per minute. The general rule is that the rate of responding is inversely proportional to the interval between reinforcements (Skinner, 1938, 1950).

Variable-interval reinforcement. The other extreme of the interval schedule is to administer the reinforcement at random times. Thus, after an animal has learned a particular response, we might reinforce it the first time after one minute, then after thirty seconds, then after three minutes, etc. Under these conditions behavior becomes extremely stable, the animal responding with almost machinelike precision.

Fixed-ratio reinforcement. In fixed-ratio reinforcement the reinforcement is regularly delivered after a specific number of responses. Here, as we might expect, the effects on performance are dramatically different from those of fixed-interval reinforcement. First of all, reinforcement on a ratio schedule places a premium on responding rapidly, for the higher the rate of responding, the higher the rate of reinforcement. If an animal is being reinforced once every 10 responses, it can get a reinforcement every 10 seconds by responding once per second. Even more characteristic of fixed-ratio reinforcement is the fact that the rate of responding is variable. Immediately after a reinforcement it is low, and it gradually increases until the next reinforcement occurs. Since this acceleration occurs as the ratio of reinforcements to responses gets low, unlike interval schedules, fixed-ratio schedules produce higher average rates of responding with fewer reinforcements.

Variable-ratio reinforcement. The variable-ratio reinforcement case is perhaps the most interesting, because it is probably more characteristic of the conditions in which natural behavior occurs. Variable-ratio reinforcement is administered after different numbers of responses. If we reinforce a well-learned habit on a variable-ratio schedule, we do not get the increase in rate of response between reinforcements obtained with fixed-ratio reinforcement since the animal cannot predict the exact number of responses necessary for a reinforcement; there is a premium on responding rapidly, for getting a reinforcement eventually depends upon the number of tries.

The effects associated with particular ways of administering reinforcement are all effects on performance—they are things that will happen to a well-learned response when a partial reinforcement schedule is introduced. Let us now look at the effects of partial reinforcement upon the specific conditions of acquisition and extinction.

EFFECTS OF PARTIAL REINFORCEMENT ON LEARNING AND EXTINCTION

The effects of partial reinforcement on acquisition. If we tried to train a rat to press a lever in the Skinner box by reinforcing it once for every 10

responses, we would have a long wait before the animal learned the response. In general this would be true in any situation in which the major thing to be learned is the relationship between a response and its consequence—the reinforcement. Since learning such a relationship is the essence of most studies on simple instrumental conditioning, it is not surprising that a review of the literature on partial reinforcement (Jenkins and Stanley, 1950) shows that most studies reveal a tendency for anything less than 100 per cent reinforcement to slow the rate of acquisition. In other words, reinforcing simple instrumental acts part of the time rather than all of the time appears to slow down the rate of learning.

This retarding effect of partial reinforcement is probably the indirect effect of the many conditions that determine an animal's performance in a given situation. We saw in the last chapter that such factors as delay of reinforcement, amount of reinforcement, etc., work primarily on performance rather than on rate of learning; any effect upon rate of learning of these reinforcement variables is indirect, determined by the fact that performance is frequently a major determiner of what animals learn. Much the same thing is true of the schedule of reinforcement; any effect it has upon learning is probably achieved through an alteration in the nature of the instrumental act or of the stimulus-response sequence that the animal must learn.

The effect of partial reinforcement on resistance to extinction. The most dramatic effect of partial reinforcement is on resistance to extinction. All schedules of reinforcement affect resistance to extinction somewhat differently, but it is a general rule that partial reinforcement will nearly always produce greater resistance to extinction than 100 per cent reinforcement (see Jenkins and Stanley, 1950).

Earlier in this chapter we saw that the number of times a particular activity had been reinforced determined, up to a limit, how many times the activity would be performed in the absence of reinforcement. The maximal resistance to extinction after a 100 per cent reinforcement schedule is not great. No matter how many times we reinforce lever pressing in the rat, it is unlikely that the animal will respond much more than 200 times without reinforcement. Thus, habits established and maintained with 100 per cent reinforcement are not very resistant to extinction.

Results obtained by Skinner (1938) on instrumental behavior and by Humphreys (1939) on a classical conditioned response showed that a schedule of partial reinforcement greatly increased resistance to extinction. Except for very special conditions, nearly every experiment since then has shown similar results. Not only does partial reinforcement make a habit

more resistant to extinction, but one experiment on maze behavior in the rat shows that such reinforcement makes it more difficult to learn a new conflicting habit (Wike, 1953). The large effects of partial reinforcement on extinction may be seen in Figure 12 (Jenkins, McFann, and Clayton, 1950).

This particular effect of partial reinforcement is of the greatest theoretical and practical importance. For one thing, it enables us to understand why behavior in natural settings is likely to be so persistent—most behavior is

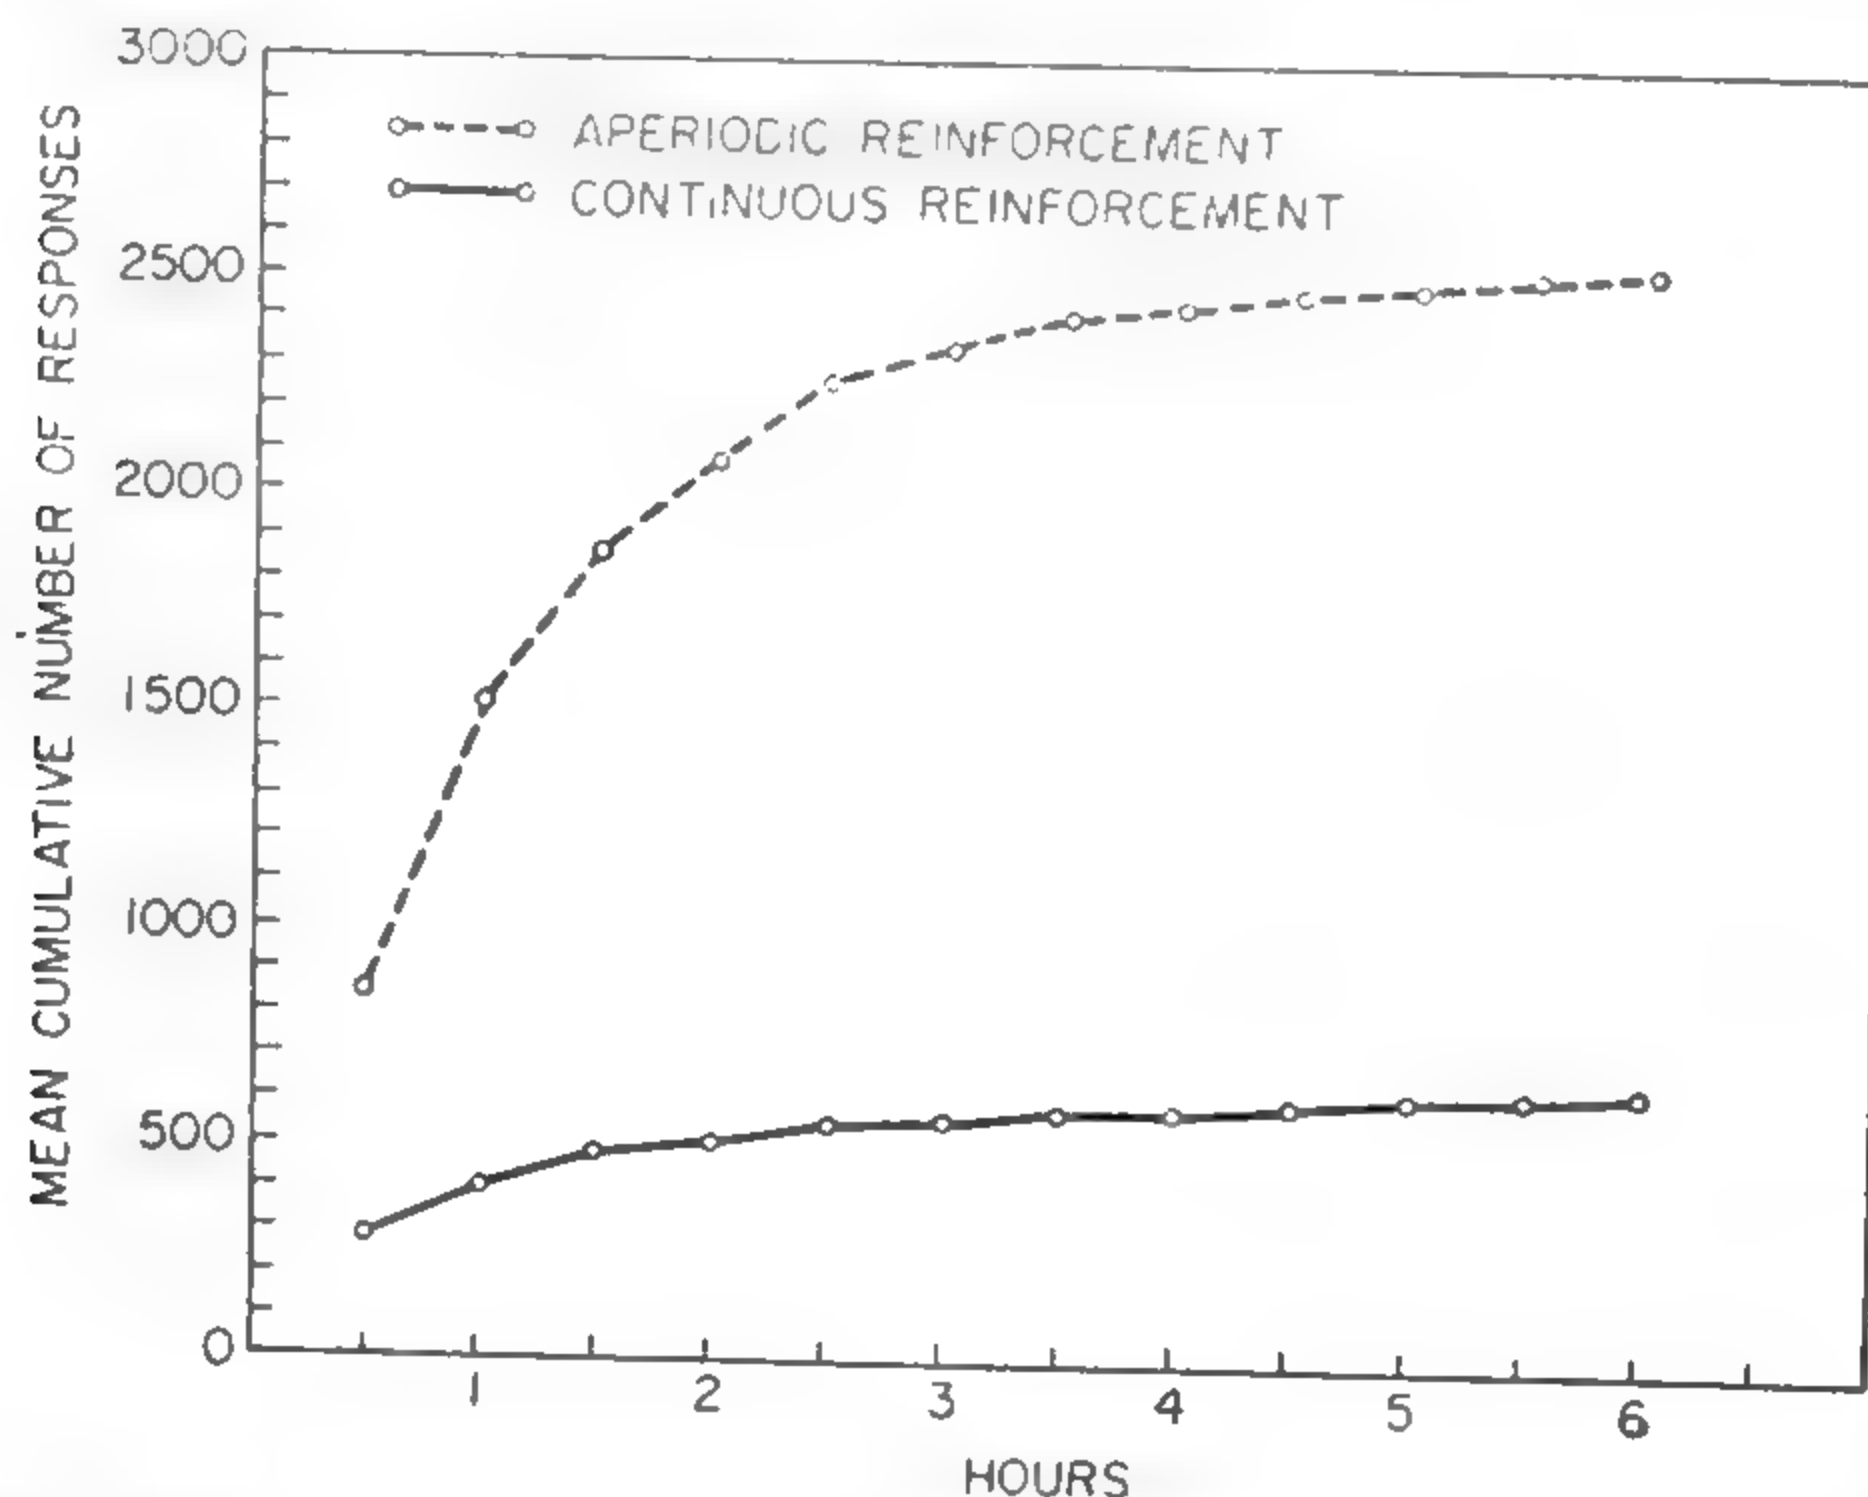


Figure 12. Cumulative extinction curves after partial and continuous reinforcement. For the partial reinforcement condition, the animals were placed on a variable-interval schedule. The animals in both groups received 200 reinforcements before extinction began. (Jenkins, McFann, and Clayton, 1950.)

probably not reinforced for each try. We also know that the most efficient use of reinforcement is obtained by a partial reinforcement schedule (for most purposes, a variable-ratio one), probably after an initial group of 100 per cent reinforced responses. In this way we conserve the strength of reinforcements that are likely to suffer from satiation effects (such as new toys do for an infant), and, at the same time, we produce the largest resistance to extinction.

There is apparently no limit to the extent to which a schedule of very infrequent reinforcement can increase resistance to extinction. Even in the laboratory it is possible to produce habits that will resist the extinguishing effects of thousands of unreinforced responses. As a matter of fact, Skinner

(1950) suggests that if conditions are proper it is possible to produce an extinction curve with no curvature; in other words, the animal goes on responding without reinforcement *ad infinitum*. Incidentally, Skinner remarks, such a result is a telling blow against a theory of extinction built upon the accumulation of a response-produced inhibition.

Some theoretical importance was originally attached to the possibility that wider spacing of trials might have the effect of reducing the heightened resistance to extinction brought about by partial reinforcement (Sheffield, 1949). Subsequent experimental work, however, has established that such a reversal does not always occur under spaced practice, and indeed probably does so only under very special circumstances (Grant, Schipper, and Ross, 1952; Weinstock, 1954).

INTERPRETATIONS OF PARTIAL REINFORCEMENT

We have seen that the importance of partial reinforcement is twofold. First of all, a comparison of different schedules of reinforcement showed that momentary characteristics of response strength depend upon the particular schedule of reinforcement under which an organism is performing. Secondly, we saw that partial reinforcement generally has the effect of increasing the resistance to extinction. This is perhaps the most fundamental point, and it is the one that has received the most attention from theorists and been applied most to theories of personality and development. Part of the interest in the effects of partial reinforcement has centered in efforts to explain why partial reinforcement so greatly prolongs the extinction process. Theories about these effects of partial reinforcement fall into two general classes—stimulus theories and response theories.

Stimulus Theories

Stimulus theories of partial reinforcement say that resistance to extinction is greater after such a schedule because the stimulus situation is different from that in 100 per cent reinforcement. Reasons for the effect of increasing resistance to extinction are spelled out in several ways. Let us look at some of these.

Expectancy. The notion of expectancy says that animals come to expect, or predict, stimulus situations associated with particular responses. Thus, a rat reinforced 100 per cent of the time for pressing a bar comes to expect food after responding. If an animal is reinforced only part of the time during training, it will not expect a reinforcement on every occasion, so that

when reinforcement stops altogether, this animal is likely to be more persistent.

This concept of expectancy has been roundly criticized on several counts. Some critics have claimed that it is an anthropomorphic notion. Others have pointed out that the definition of expectancy is vague and ambiguous. Certainly, the criticism of anthropomorphism is justified. Human learners are capable of expectancy—as a form of verbal behavior. We can tell ourselves when we ought to expect something to happen. It is doubtful that animals are capable of expectancy in this sense.

The most serious criticism of the notion of expectancy is that it is an *ad hoc* explanation. It is acceptable enough in itself if, for example, it is defined as the ability to respond to a conditioned cue before the unconditioned stimulus appears. When the buzzer sounds the dog salivates; then the meat powder is placed in its mouth. If this is all one means by expectancy, there is no quarrel. This definition of expectancy (and it is nothing more than a definition) does not provide us, however, with a very powerful explanatory tool. We may apply it to the problem of partial reinforcement, and by making some further incidental definitions, come out with the prediction that partial reinforcement will lead to greater resistance to extinction than 100 per cent reinforcement.

We have dealt rather harshly with this idea of expectancy. Later in this book we shall examine questions about decision making in human beings in which the notion of expectancy has been considerably more useful than it has been in the theory of simple instrumental and classical conditioned responses.

Discrimination. Another simple idea used to account for the results of partial reinforcement studies states that under partial reinforcement schedules it is less easy for the animal to discriminate between conditions of reinforcement and those of extinction. This implies that the major determiner of the rate (and probably the limit) of responding during extinction is the extent to which the situation during extinction is the same as that during original reinforced training. In partial reinforcement there are occasions when responses go unreinforced; consequently, when extinction is eventually begun it is more difficult for the animal to discriminate between the conditions of reinforcement and those of extinction.

It would seem that this discrimination notion is quite similar to a general concept of expectancy. The difference is that the discriminability hypothesis is stated in such a way as to preclude our attributing symbolic, human processes to rats as they generate extinction curves. The word expectancy has

too much excess meaning for most people. Perhaps it is the better part of scientific caution to accept the discriminability notion as part of the explanation for the partial reinforcement effect and to dismiss the concept of expectancy.

The notion of discriminability has other advantages, as we shall see later. The principal of these is that we know some of the rules for discrimination learning (see Chapter 4) and can use them to predict the specific effects of particular patterns of partial reinforcement. At any rate we can accept the notion that the increased resistance to extinction resulting from partial reinforcement can be considered as partly due to a change in the discriminability between the conditions of reinforcement and those of extinction.

Response Theories

Another way to look at the question of what produces the increased resistance to extinction with partial reinforcement is to consider the definition of the response unit. For most purposes we define the unit of response as that which will produce a reinforcement. Thus, in the Skinner box, it is a lever push capable of operating the feeder. It does not make any difference how the lever is pressed, as long as it is pressed hard enough. By the same token we might argue that in any partial reinforcement situation the unit response is defined by the *number* of lever presses necessary to activate the feeder. If we did this, we should consider all of the unreinforced lever pushes between reinforced ones to be part of the activity necessary for reinforcement and count them as only one response. Thus, if on a variable-ratio reinforcement schedule we reinforced 1 lever press in 10, we would count 10 lever presses as only 1 response. On this variable-ratio reinforcement schedule 10 lever presses would be equivalent to 1 on a 100 per cent reinforcement schedule.

This argument was advanced by Mowrer and Jones (1945), and they showed that when responses are counted by this scheme, actually *fewer* responses occur during extinction after partial reinforcement, though, of course, more single bar presses occur. The reason that fewer responses occur, argued Mowrer and Jones, is that a stronger response-induced inhibition is induced by the greater activity required under partial reinforcement. It would have been interesting to see if this were also accompanied by greater spontaneous recovery, but Mowrer and Jones did not examine this question.

While we have presented this notion as an effect upon responding and the hypotheses of expectancy and discriminability as effects upon stimuli, this

is, in the last analysis, largely a matter of convenience. Furthermore, there is nothing basically incompatible between these notions as they stand. It is possible that the unit of response as far as the organism is concerned is determined by what it "expects" or by the occurrence of a reinforcement. At any rate, while a more adequate explanation of the effects of partial reinforcement upon extinction may eventually be stated, it is probable that it will be close to those discussed above.

Importance of partial reinforcement. The effects of partial reinforcement—particularly on extinction—are extremely important when applying the principles of learning and behavior theory to the understanding of behavior under natural conditions. Most habits seem to be highly resistant to the effects of extinction, and this is very likely because behavior in natural life is reinforced on a probable rather than a certain basis. Thus, the youngster who is beginning to get a command of language is likely to be persistent about asking for cookies, ice cream, etc., because, in the course of things, he may be reinforced on a variable-ratio schedule. Sometimes these variable-ratio reinforcements produce behavior which is highly resistant to the effects of extinction. Parents who are concerned about controlling the behavior of their offspring will sometimes try to introduce fixed-interval reinforcement. Thus, some children are supposed to be able to buy ice cream on Mondays and Thursdays but not on Tuesdays, Wednesdays, and Fridays. Since parents are not nearly so consistent as the schedules of reinforcements set up for rats in Skinner boxes, the discriminations based on fixed-interval reinforcement programs that appear in laboratory rats do not always appear to be clear-cut in children.

Thus, inconsistency in the application of reinforcement is the stuff of which persistent behavior is made. How much an inconsistent application of reinforcements (and punishments) is responsible for a generalized trait of persistence we do not know. Not very much, it would seem, for even though parents differ enormously in the degree to which reinforcement and punishment are consistently applied, the rest of the world is likely to be much the same in this respect. Each youngster, at least in the easygoing atmosphere of American culture, is likely to have been exposed to a fairly consistent pattern of inconsistency in the application of reinforcement.

Partial reinforcement in choice and decision. A problem that has received considerable attention recently is that of finding out how irregular patterns of reinforcement affect decisions. More generally, this has been conceived of as a problem in the probabilistic nature of events that influence human and animal behavior. Brunswik (1954) for many years insisted that

the major variables in psychology were to be considered as only probabilistic, not certain, in nature. For instance, a cue for visual depth, or size, of retinal image, does not always yield unequivocally correct information. Objects that loom large in our visual field are *probably* objects close to us in space, but not necessarily so. So it is with reinforcements; they occur in a probabilistic way, seldom in an all-or-nothing way.

With this general view in mind, Brunswik (1939) reported an experiment on rats in a T maze in which most of the time reinforcements occurred on the right side but sometimes occurred on the left. Under these conditions, rats came to distribute their choices according to the probability of reinforcement. Thus, if the right side were reinforced 75 per cent of the time, and the left side only 25 per cent of the time, the rats would choose to go to the right side 75 per cent of the time.

Of great current interest in the theory of behavior is what determines the likelihood of an organism making a particular choice given the probability but not the certainty of a particular event happening. Most experimental research has been directed toward verbal choice in adult human beings. In a typical experimental situation, a subject is presented with a panel on which are two lights and two buttons. He is to predict by pressing a button which of the lights is going to flash on a particular trial. The experimenter has arranged a particular sequence of order for flashes on the right or left. As with Brunswik's rats in the T maze, human beings apparently choose the alternatives in proportion to the frequency of the to-be-predicted event happening (Grant, Hake, and Hornscth, 1951; Estes and Straughan, 1954, etc.). This is a puzzling situation, for it is not the optimal strategy. We are not surprised, perhaps, that rats do not maximize the probability of receiving a reinforcement, but we are perhaps a little surprised that college students do not maximize the probability of being right (given a situation in which the frequency of alternative events is not even the best strategy is always to bet on the more frequently occurring event, if we have no further information).

It has been argued that this is a consequence of the learning mechanisms (Estes and Straughan, 1954; Bush and Mosteller, 1955), but if this is so it is an oversimplification. Hake (1955) has reviewed much of the literature on this topic and concluded that the conditions determining probability of a human subject making a given choice between two alternatives are influenced by several variables. There are reasons for believing, he points out, that in any such choice the subject is as much influenced by short sequences of events just before a particular choice as he is by the over-all frequency

of past events. Also, to some extent, subjects are influenced by the nature of their previous choices. Hake points out that the matching of objective probability is a result; it does not necessarily describe what individuals do.

Is this true of Brunswik's rats? It is perhaps too early to arrive at a conclusion. There are superficial parallels between the behavior of nonverbal subjects such as rats, pigeons, and young children under partial reinforcement, and the behavior of human adults in choice games with varying probabilities of making the right guess. Beyond this we can say little; we do not know if the parallels are anything more than superficial, and it is too early even to guess. We do know, however, that these experiments on human choice behavior that resemble certain experiments on partial reinforcement in animals so much are looking into important questions, and may eventually have serious implications for all matters that involve human beings making decisions or choices.

CHAPTER 4

DISCRIMINATION AND LEARNING

This
this

This
This

The chief function of learning for most animals in the world is to enable them to emit the right response at the right time. This generally means performing a response, on signal, that is already a part of the basic behavior of the animal. Most organisms seldom if ever find it necessary to learn skilled acts; much of the behavior (considered purely as response) of animals is unlearned. They must frequently learn, however, under what circumstances certain behavior is appropriate. It is the circumstances of emission rather than the behavior per se that require learning in most animals. To be sure, in the laboratory we sometimes require rats to perform skilled acts, such as dropping marbles down a narrow chute, but this kind of thing is usually highly unnatural. More typically we take advantage of something a rat does more readily, such as exploring a narrow alley. Consequently a good deal of the learning a rat does in the laboratory involves simply the discrimination of stimuli in the environment and attaching appropriate responses to these stimuli. Some psychologists have said that in this sense all learning is discrimination.

Much of innate behavior, particularly in the lower organisms, is highly discriminative; some instinctive act is likely to be emitted only in the presence of a rather precisely defined (unconditioned) stimulus. This has always been more or less understood, but in recent years we have greatly extended our knowledge of the precision with which particular stimuli set off patterns of innate action. For example, in the three-spined Stickleback fish, a very crude model of a female will elicit the innate pattern of mating behavior in the male, providing the model is painted red on its underneath surface (Tinbergen, 1951). Here the critical feature of the stimulus seems to be the red color.

In this chapter we shall look at the way in which animals learn to dis-

criminate. Since our main purpose is to understand the nature of learning—particularly as it applies to the human animal—we shall ignore some fascinating accounts of unusual discriminations in nature in favor of a detailed treatment of the process of learning to discriminate. We should note that this is not a process that is limited to learned stimuli; some unconditioned stimuli have their effects modified through discriminative training (see Tinbergen, 1953).

Before we look at discrimination itself, however, let us take a brief look at its opposite, generalization. Here we ask, How much variation can we introduce into a response-producing stimulus before the response disappears?

GENERALIZATION

While investigating the conditioned salivary response in dogs, Pavlov soon discovered that the conditioned response was not limited to the stimulus originally conditioned. It would spread to other, related stimuli. Experiments from Pavlov's laboratory (Pavlov, 1927) established the fact that if tactile stimulation of a small area of a dog's skin is used as a conditioned stimulus, stimulation of other areas of the skin will also elicit the conditioned response, though the effect diminishes with the distance from the original conditioned stimulus. This important effect is what is now known as *stimulus generalization*.¹ The two important points about stimulus generalization are (1) that the effectiveness of conditioning is not limited to the stimulus originally conditioned, and (2) that the ability of a stimulus to elicit a conditioned response decreases with distance from the original conditioned stimulus. Since both of these points are basic to the psychology of learning, let us look at some experimental evidence about them.

Generalization of classical conditioned responses. The experiments from Pavlov's laboratory are, of course, experiments on the generalization of classical conditioned responses. Let us, however, look at some experiments from laboratories that have used other classical conditioned responses and have more carefully controlled the quantitative relations between stimuli.

In one experiment, the investigators conditioned the galvanic skin response² in man (elicited by an electric shock) to a vibrator applied near the

¹ "Irradiation" and "induction" are some of the other terms that have been used to describe stimulus generalization.

² The galvanic skin response (GSR) is the electrical activity of the sweat gland cells when they are excited. These glands are largely innervated by the sympathetic branch of the autonomic nervous system; hence this response is an "emotional" one.

shoulders (Bass and Hull, 1934). The strength of the conditioned response was then tested with various stimuli placed at increasing distances from the original conditioned stimulus. The results of this experiment are shown in Figure 13. The conditioned galvanic skin reaction was greatest when it was nearest the original conditioned stimulus, and diminished in size as the stimuli increased in distance from the original conditioning site.

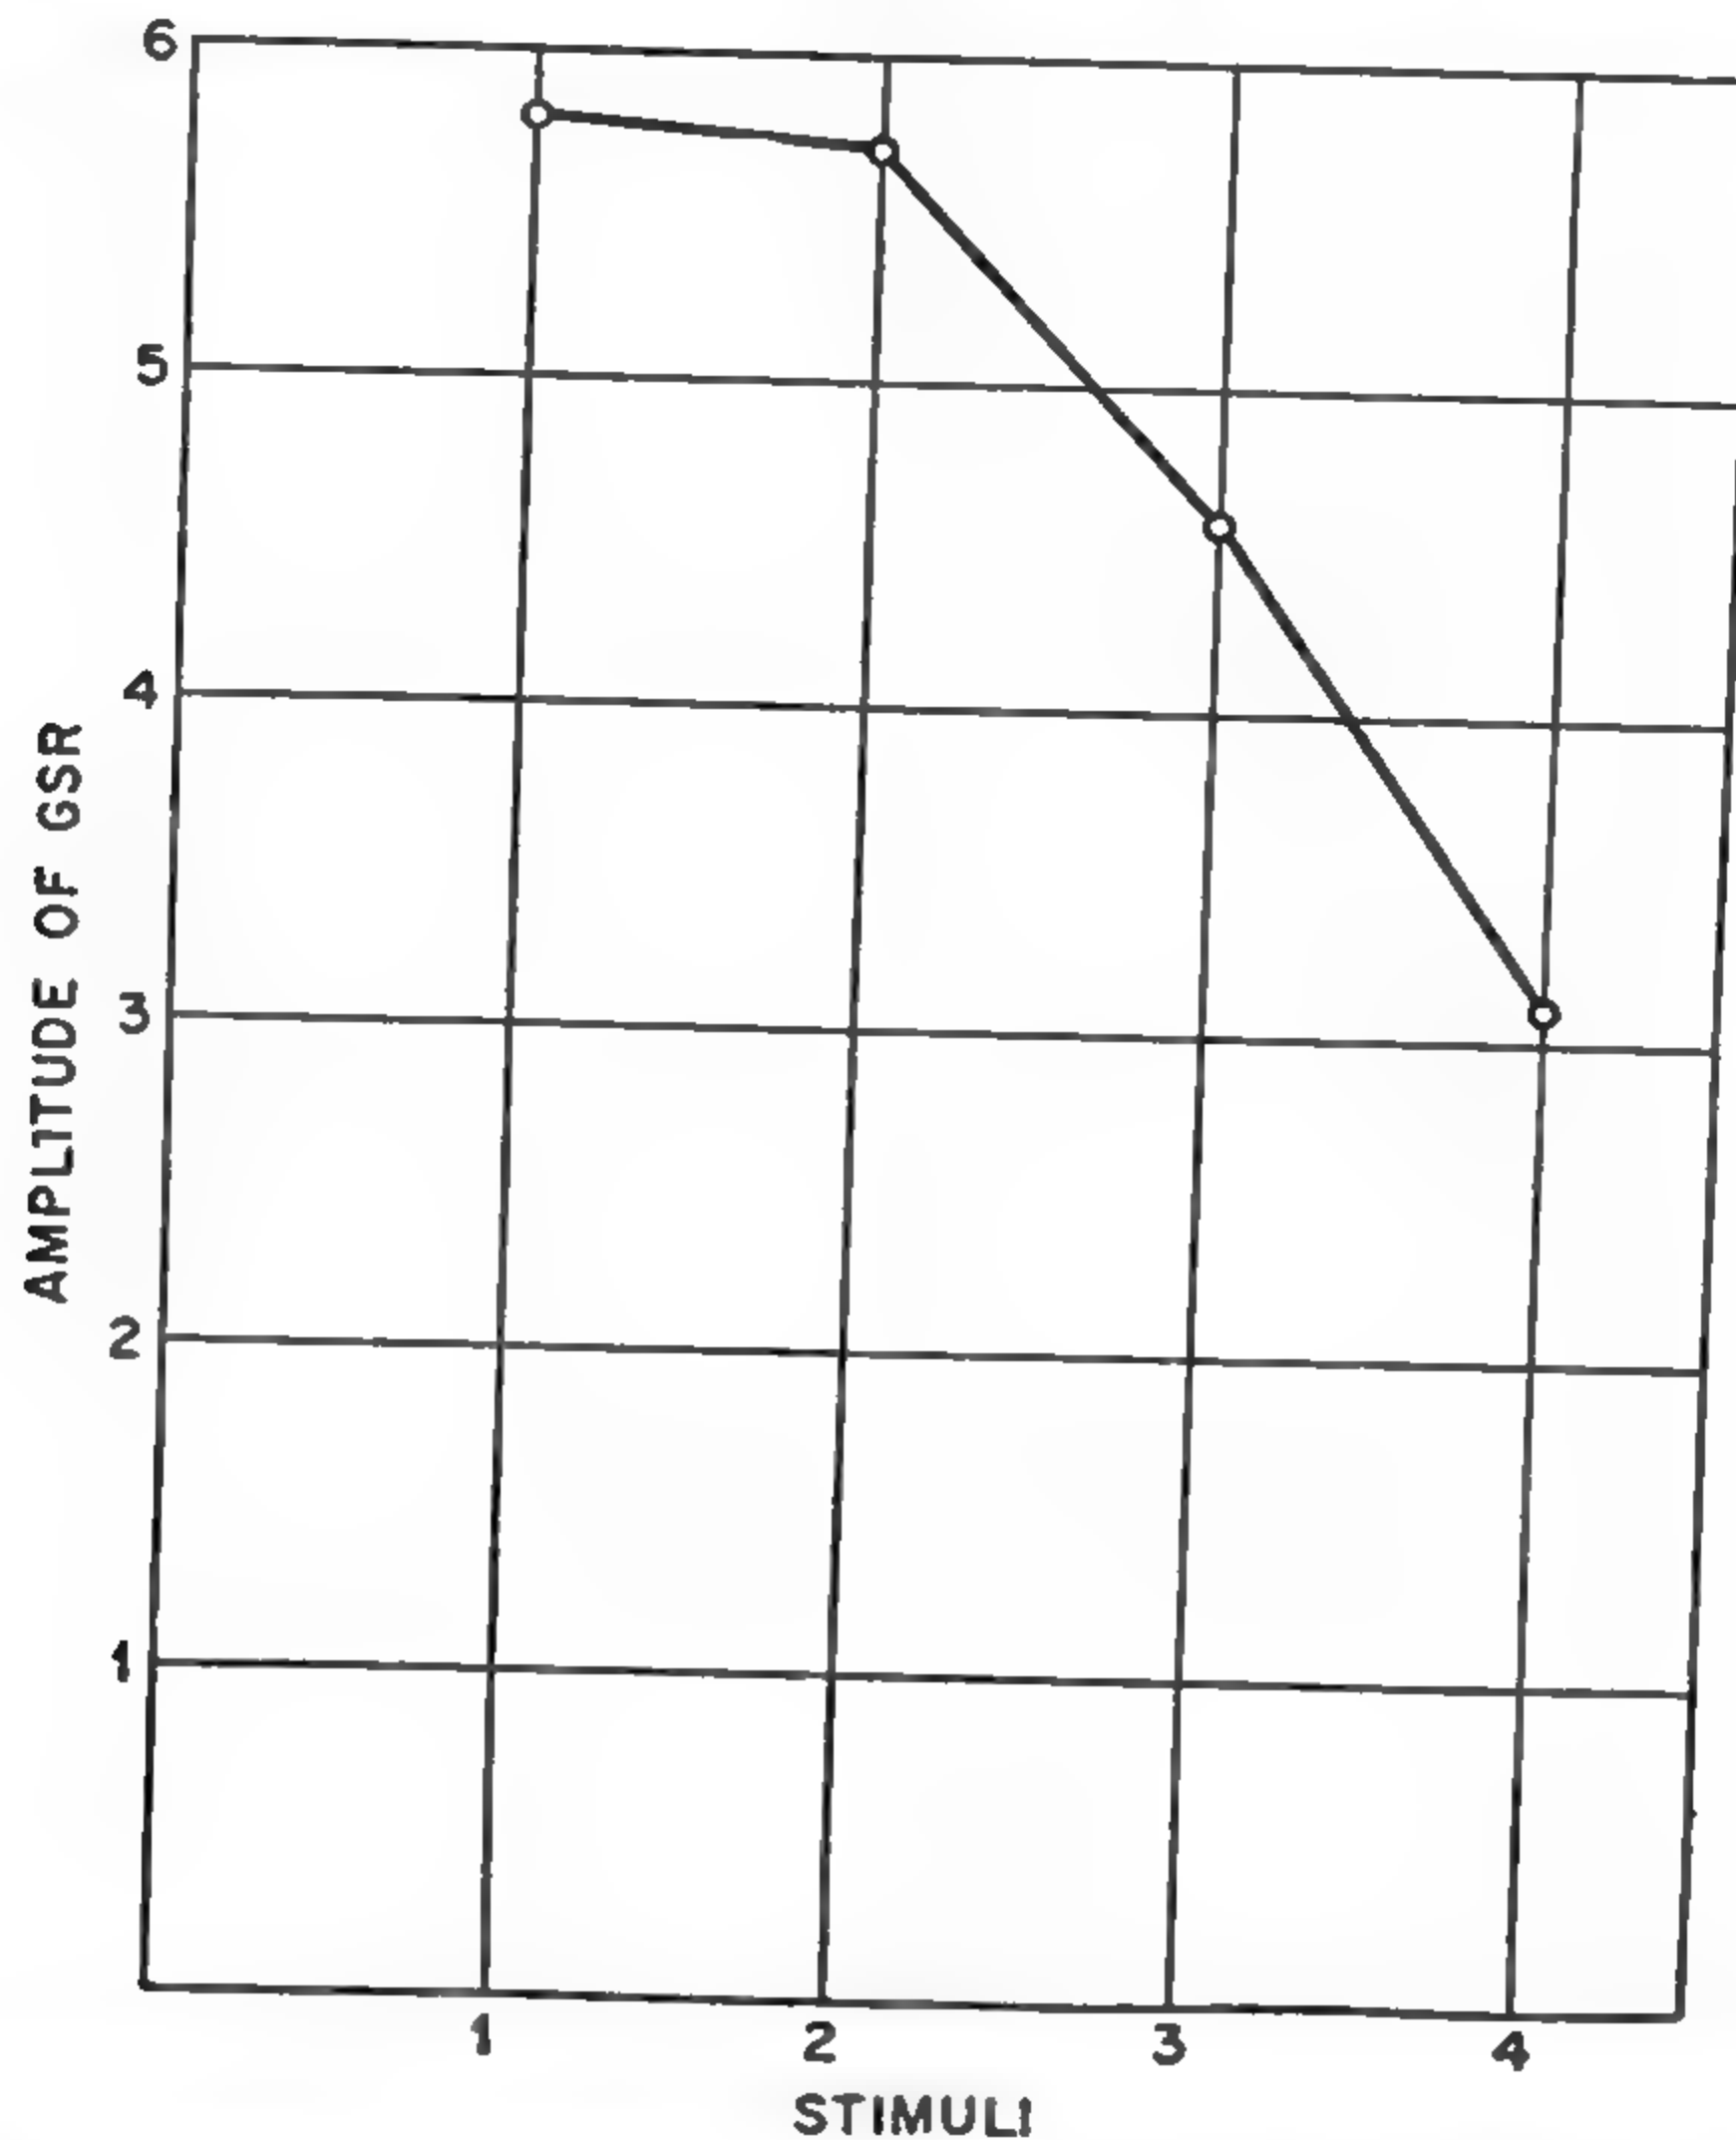


Figure 13. Generalization of a conditioned galvanic skin response to different stimuli. Stimulus 1 is the conditioned stimulus used during training. All stimuli used for testing differed from the conditioned stimulus only in location. (Data from Bass and Hull, 1934.)

There have been many experiments on the generalization of the conditioned galvanic skin response to a variety of stimuli—auditory, tactual, and visual (Hovland, 1937a, 1937b; Littman, 1949; Grant and Schiller, 1953; etc.). Most of these studies were not performed with the idea of establishing the validity of the principle of generalization—there seems to be little doubt about that—but they were attempts to establish precisely the nature of the *generalization gradient*. The reason for all this experimental activity is that there has been a widely accepted notion that there must be some

basic mathematical law governing the amount of stimulus generalization (Hull, 1950). Therefore a large number of experiments have been performed in an effort to discover the nature of this law. The sad outcome of all of these experiments suggests that there is not any such law. Rather, empirically determined gradients of generalization take a variety of forms and have many different slopes. Some of the possibilities for these are illustrated in Figure 14.

Actually there are good reasons why an exact mathematical law of stimulus generalization has not been found. Some of these are purely mechanical. For example, even for the case of the galvanic skin response, there are

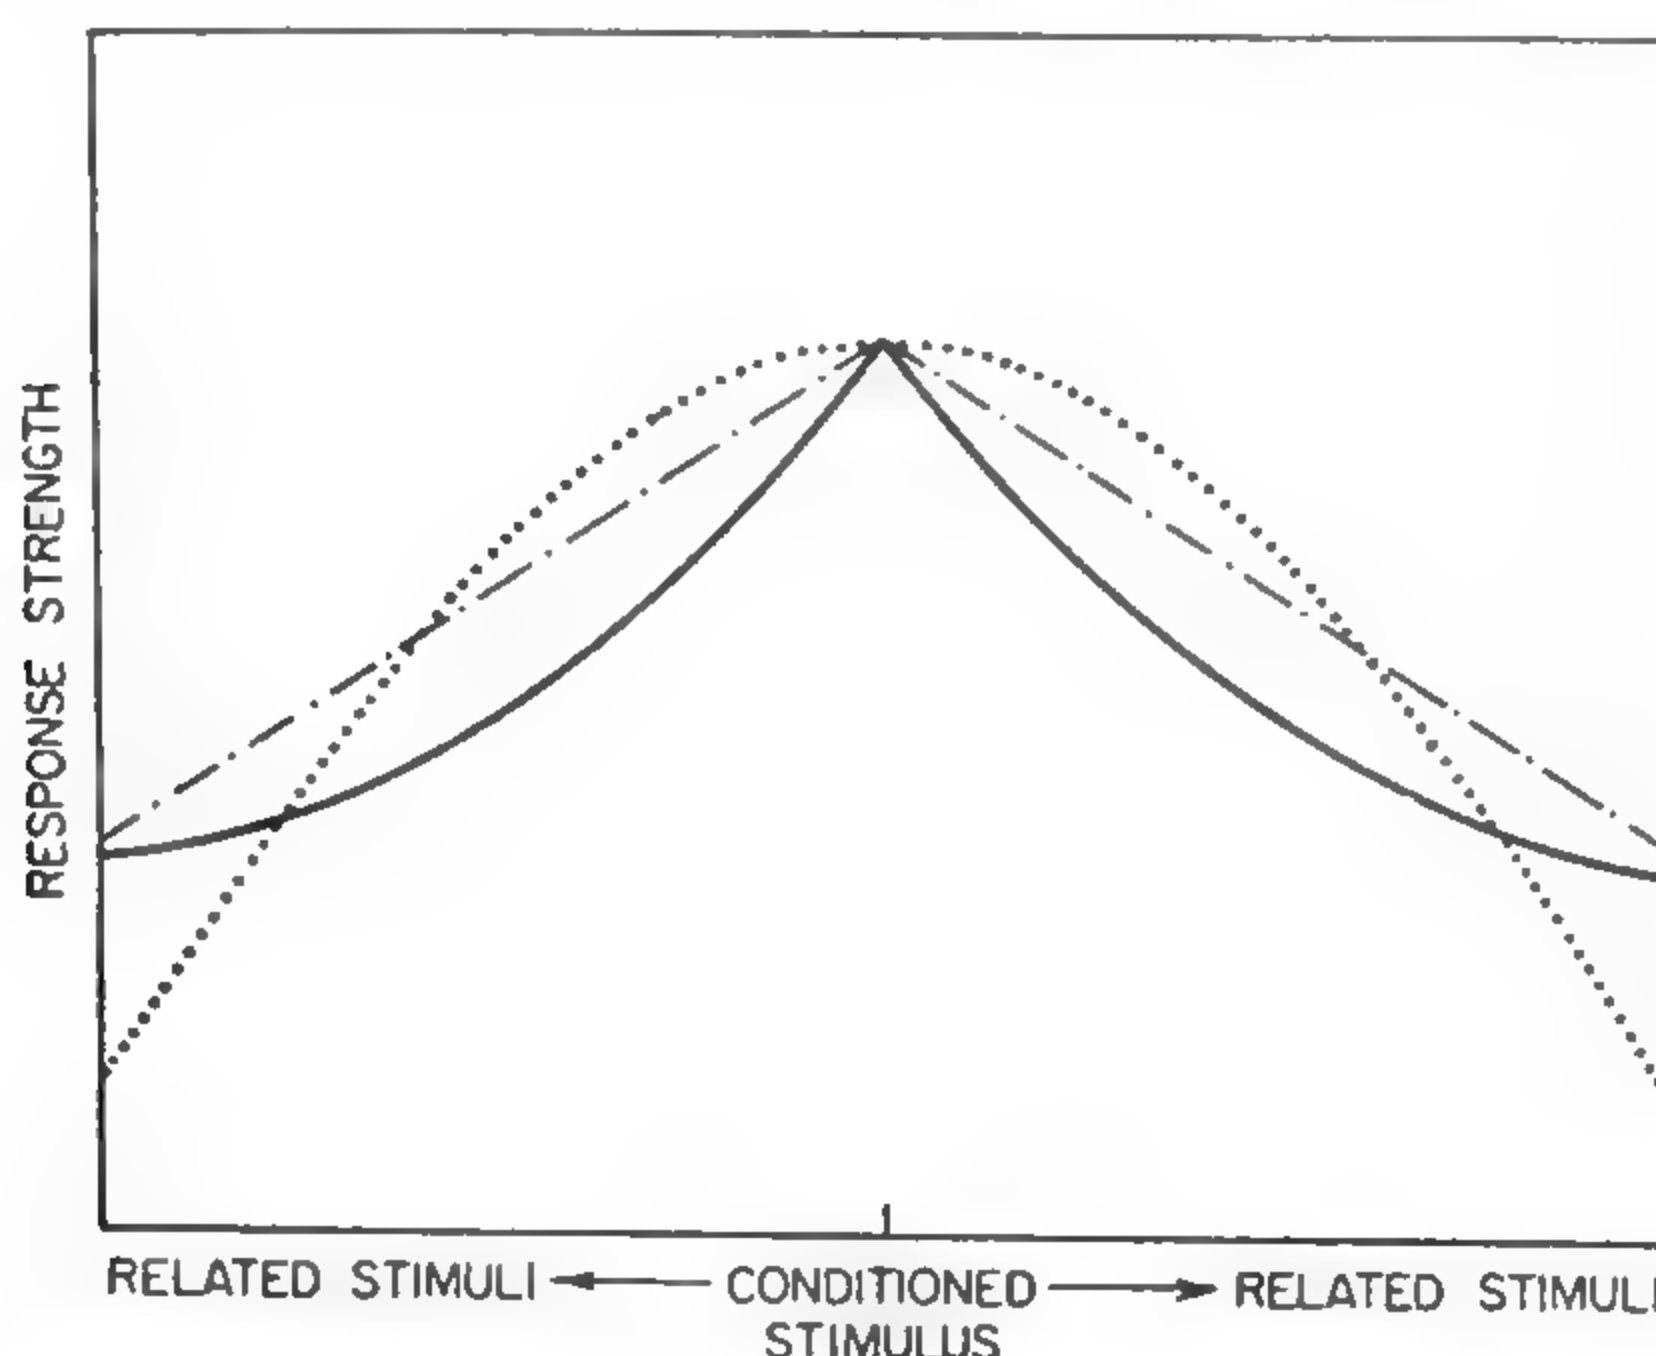


Figure 14. Some theoretical possibilities for generalization curves. Each curve would be expressed by a different mathematical equation.

many ways in which the response can be measured—as electrical resistance, as current flow, as electrical potential, and, if a-c circuits are used, as other quantities. Unfortunately there is no rational choice for the basic measurement that should be employed. This precludes a fundamental mathematical law of stimulus generalization for even so restricted a case as the galvanic skin response. When we try to generalize a mathematical law to other responses, salivation, for example, the situation becomes even more hopeless, for here we introduce other arbitrary ways of measuring.

Some theorists maintain that if probability of response occurrence is used as the sole psychologically meaningful measure of behavior, we can then establish a mathematical law of generalization. The difficulty is that there is no unequivocal and psychologically meaningful way of stating the

units on the sensory dimension along which generalization is supposed to occur. Bush and Mosteller (1951) do suggest an index of similarity between stimuli, but they are unable to provide any way of rationally relating this index to the physical characteristics, such as size, distance, and intensity, that we use in measuring stimuli. Thus they are unable to suggest an exact mathematical statement of the principle of stimulus generalization. It is possible that this problem will soon be solved, but at present we must be content with the qualitative principles that stimuli do generalize their conditioned effects and that they generalize roughly to the extent that other stimuli are related to or resemble them.

Generalization of instrumental conditioned responses. Generalization in classical conditioning is a relatively straightforward matter, since in classical conditioning some stimulus is always used to elicit the conditioned response. In instrumental conditioning, however, we do not always attach a response to some particular stimulus, though it is usually perfectly feasible to do so. We can, for example, train a rat to press a lever in the Skinner box only upon signal. Such training ordinarily involves some discriminative learning, and generalization which comes after discrimination learning is quite different from that which comes before. Consequently in most instrumental learning we look at generalization only after some discriminative learning has gone on; in order, however, to make a more exact parallel with classical conditioning, let us examine a case of stimulus generalization in instrumental conditioning in which there has been little or no discrimination training.

In this study (Grice and Saltz, 1950) rats were trained to run down an alley and obtain food by pushing back a small door. The door was embedded in the stimulus to which the rats were trained; it consisted of a white circle on a black background. The rats were trained to run to a stimulus of a certain size, say, 79 square centimeters, and then were tested with circles of different sizes (for example, 50 square centimeters). The test consisted in finding out how many times the rats would run during extinction (no reinforcement).

The results of the experiment are shown in Figure 15. As we can see, well-defined gradients of generalization occurred. The more remote the size of the test stimulus from the training stimulus, the fewer the responses during extinction. Thus it is clear that if a reinforcement is associated with a particular stimulus during instrumental training, some of that association will "spill over" to stimuli related to, but not the same as, the stimulus originally conditioned.

The nature of generalization. The existence of gradients of stimulus generalization both for classical and instrumental conditioned responses is not in doubt. There has been, however, considerable argument over the nature of these gradients. Lashley and Wade (1946), for example, have argued that they simply represent absence of opportunities for discrimination. In other words, they argue, if animals are not allowed to compare stimuli simultaneously, they find it difficult to discriminate between closely related

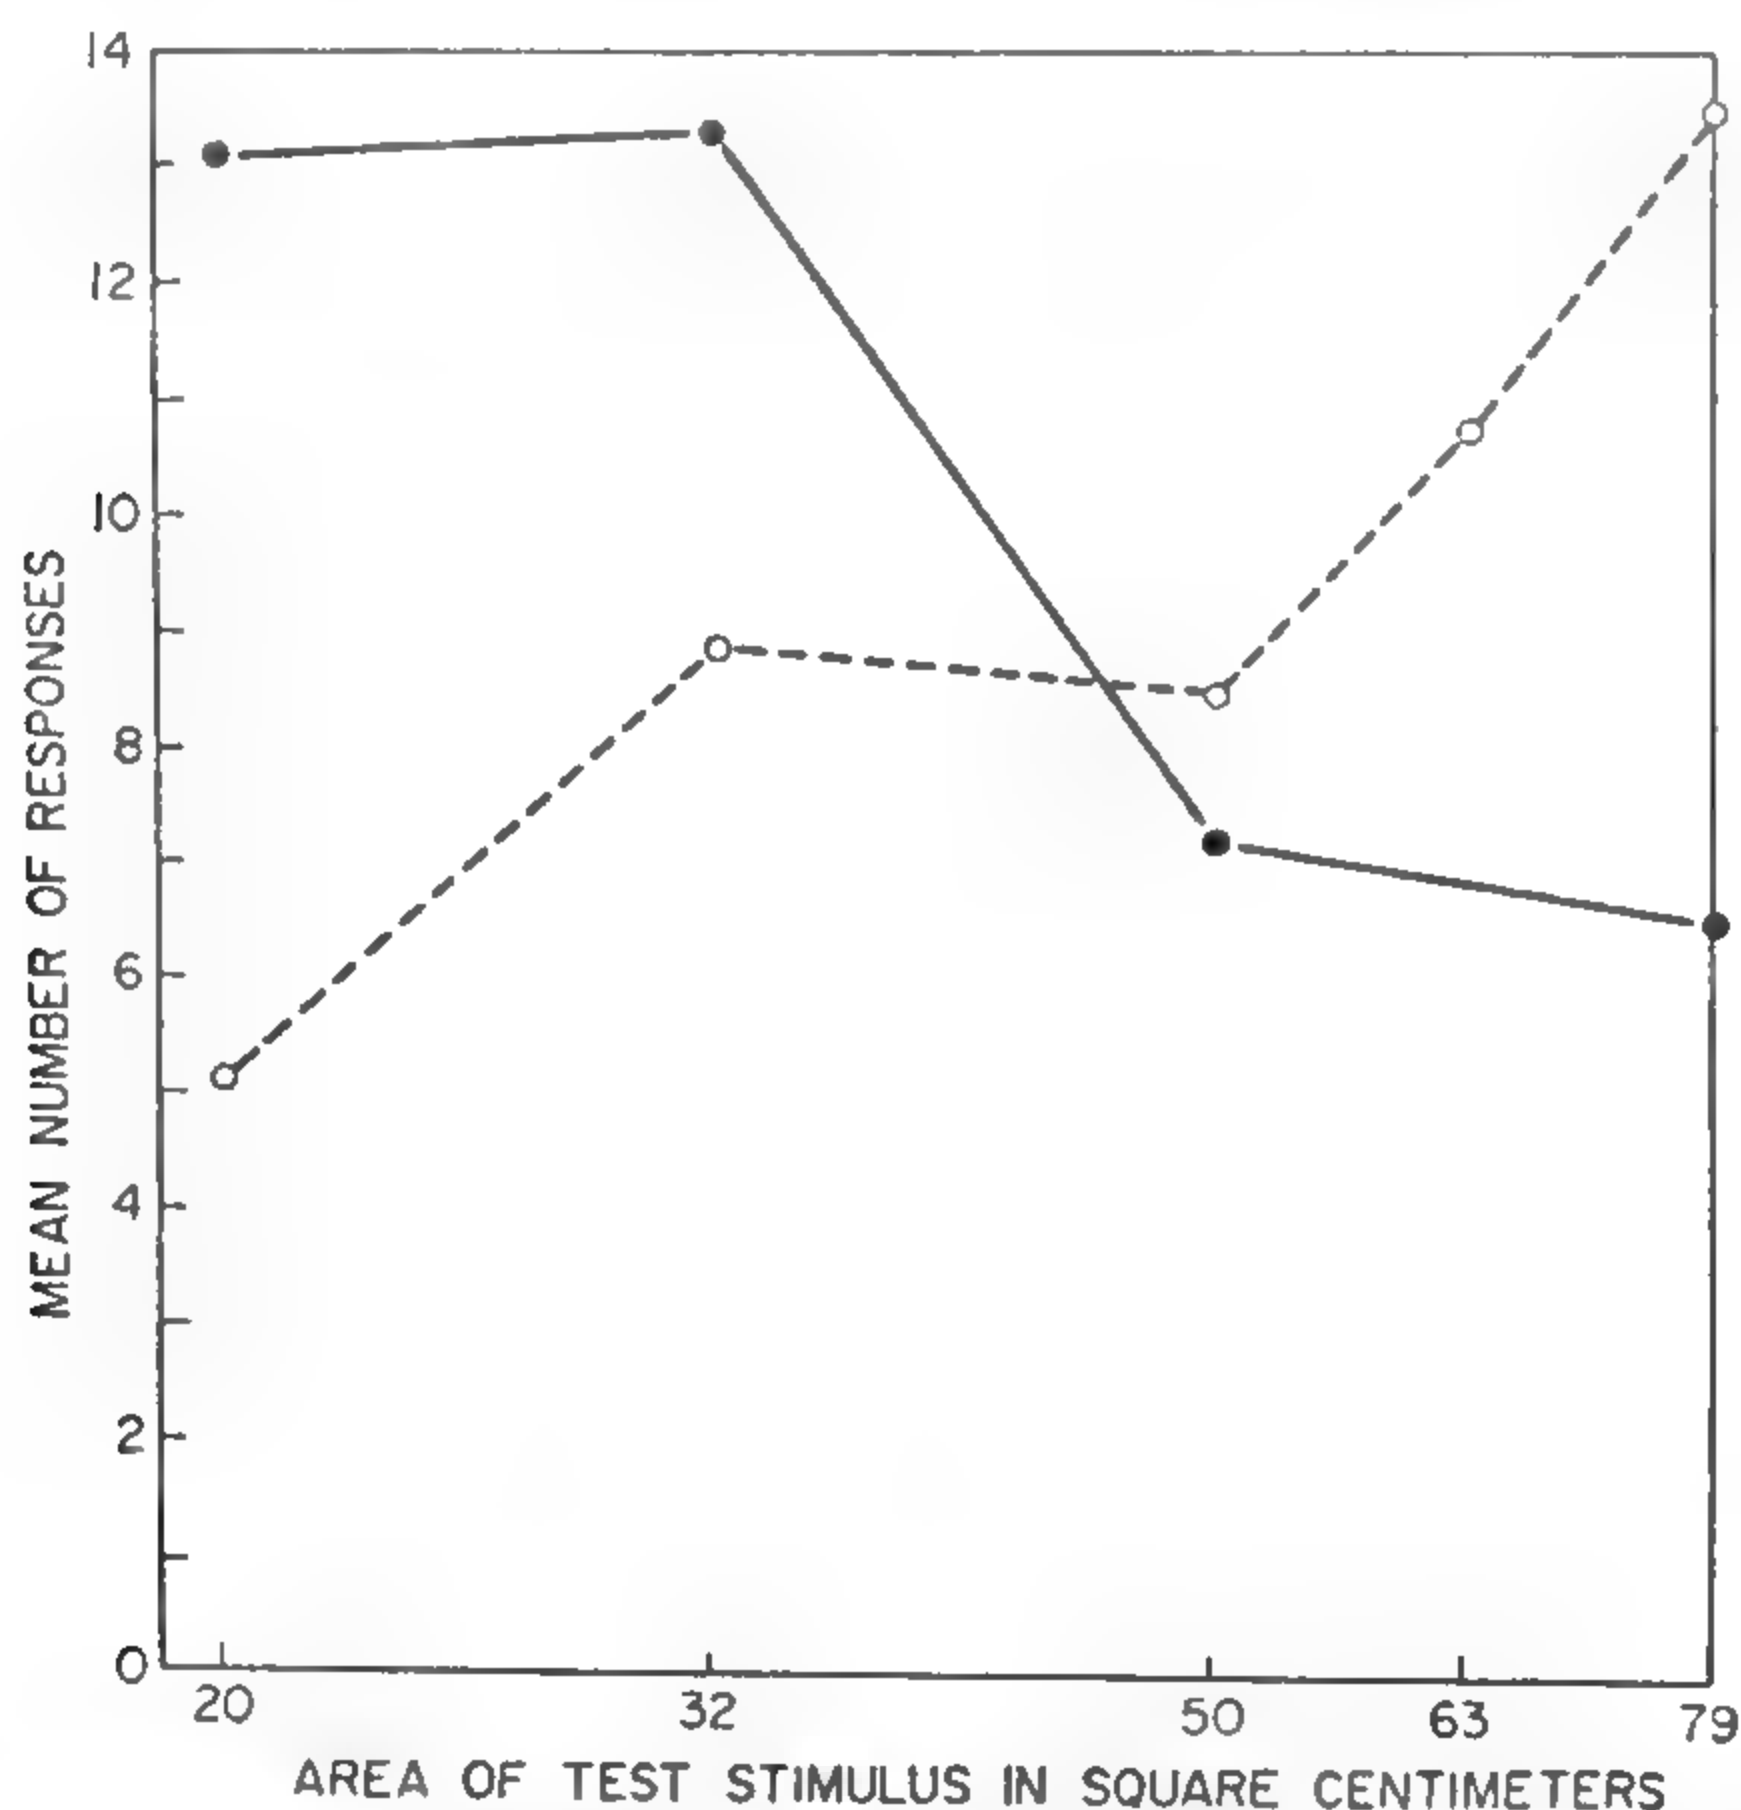


Figure 15. Generalization gradients for an instrumental response. The two curves show the mean number of responses during extinction to various test stimuli. The solid curve shows generalization after training on the 20 square millimeter stimulus, and the dotted curve shows the same thing on the 79 square millimeter stimulus. (Grice and Saltz, 1950.)

stimuli. Thus, generalization does not represent a spill-over of the influence of habit so much as it does the failure to discriminate along some sensory dimension. The fact is, as we shall see later, that animals can learn to discriminate stimuli even when they are presented successively rather than simultaneously, so any failure of discrimination in experiments on generalization more than likely represents a failure of *learned* discrimination, not a failure in sensory capacity for discrimination.

It is probable that the particular generalization gradient exhibited by an organism under test in the laboratory, even if no specific discriminative

training has been given, represents the effects of the history of an organism with sensory learning. Even a laboratory rat that has been caged in an uneventful animal colony all its life has had some opportunity for discriminative learning, and this probably has some effect on laboratory tests. We do not know much about these effects, but ever since the appearance of Hebb's book *The Organization of Behavior* (1949), learning theorists have become increasingly aware of their existence.

Extinction also generalizes. We have seen that the effects of reinforcement generalize to stimuli other than those specifically used in training. As might be expected, the effects of extinction also generalize in this way. Let us illustrate this with an example from classical conditioning (Hovland, 1937a). Suppose we condition a galvanic skin response in human subjects by pairing an electric shock with a tone of a certain frequency. After conditioning we shall, of course, find that another frequency will elicit the response. Suppose we then extinguish the conditioned response to the *second* tone. If we then go back and test with the first tone, we shall find that extinction of the response to the second tone also reduced the tendency to respond to the first. Thus the effects of extinction generalize as do the effects of reinforcement.

The generalization of the effects of both reinforcement and extinction are extremely important in the theory of discriminative learning, and a large portion of this chapter, as well as later chapters, will be devoted to the application of the principles of stimulus generalization to an understanding of discriminative learning and of related questions in the psychology of learning.

DISCRIMINATIVE LEARNING

Learning to make discriminations can be looked upon as the process of breaking down generalizations. This may be achieved by reinforcing responses to one stimulus and extinguishing responses to another. Discrimination is achieved when the organism comes to respond to one stimulus and not to the other.

There are a variety of specific procedures for producing discriminative learning in laboratory animals. One of the best known of these involves the use of the Lashley jumping stand. An example of this device is shown in Figure 16. In studies of discriminative learning using the jumping stand, rats are trained to jump from a platform about 8 inches from the stand to one of two doors in the stand. Light pieces of cardboard are placed over the

doors, and these carry the stimuli. The problem illustrated in Figure 16 requires the rat to discriminate between vertical and horizontal stripes. If the rat makes the correct choice, the door falls and the rat is allowed to eat food placed inside the stand. If it makes the wrong choice, the door does not move and the rat falls to the net below. Thus, typically, the Lashley jumping stand makes use of punishment.

In other types of discriminative studies single choice-point mazes and

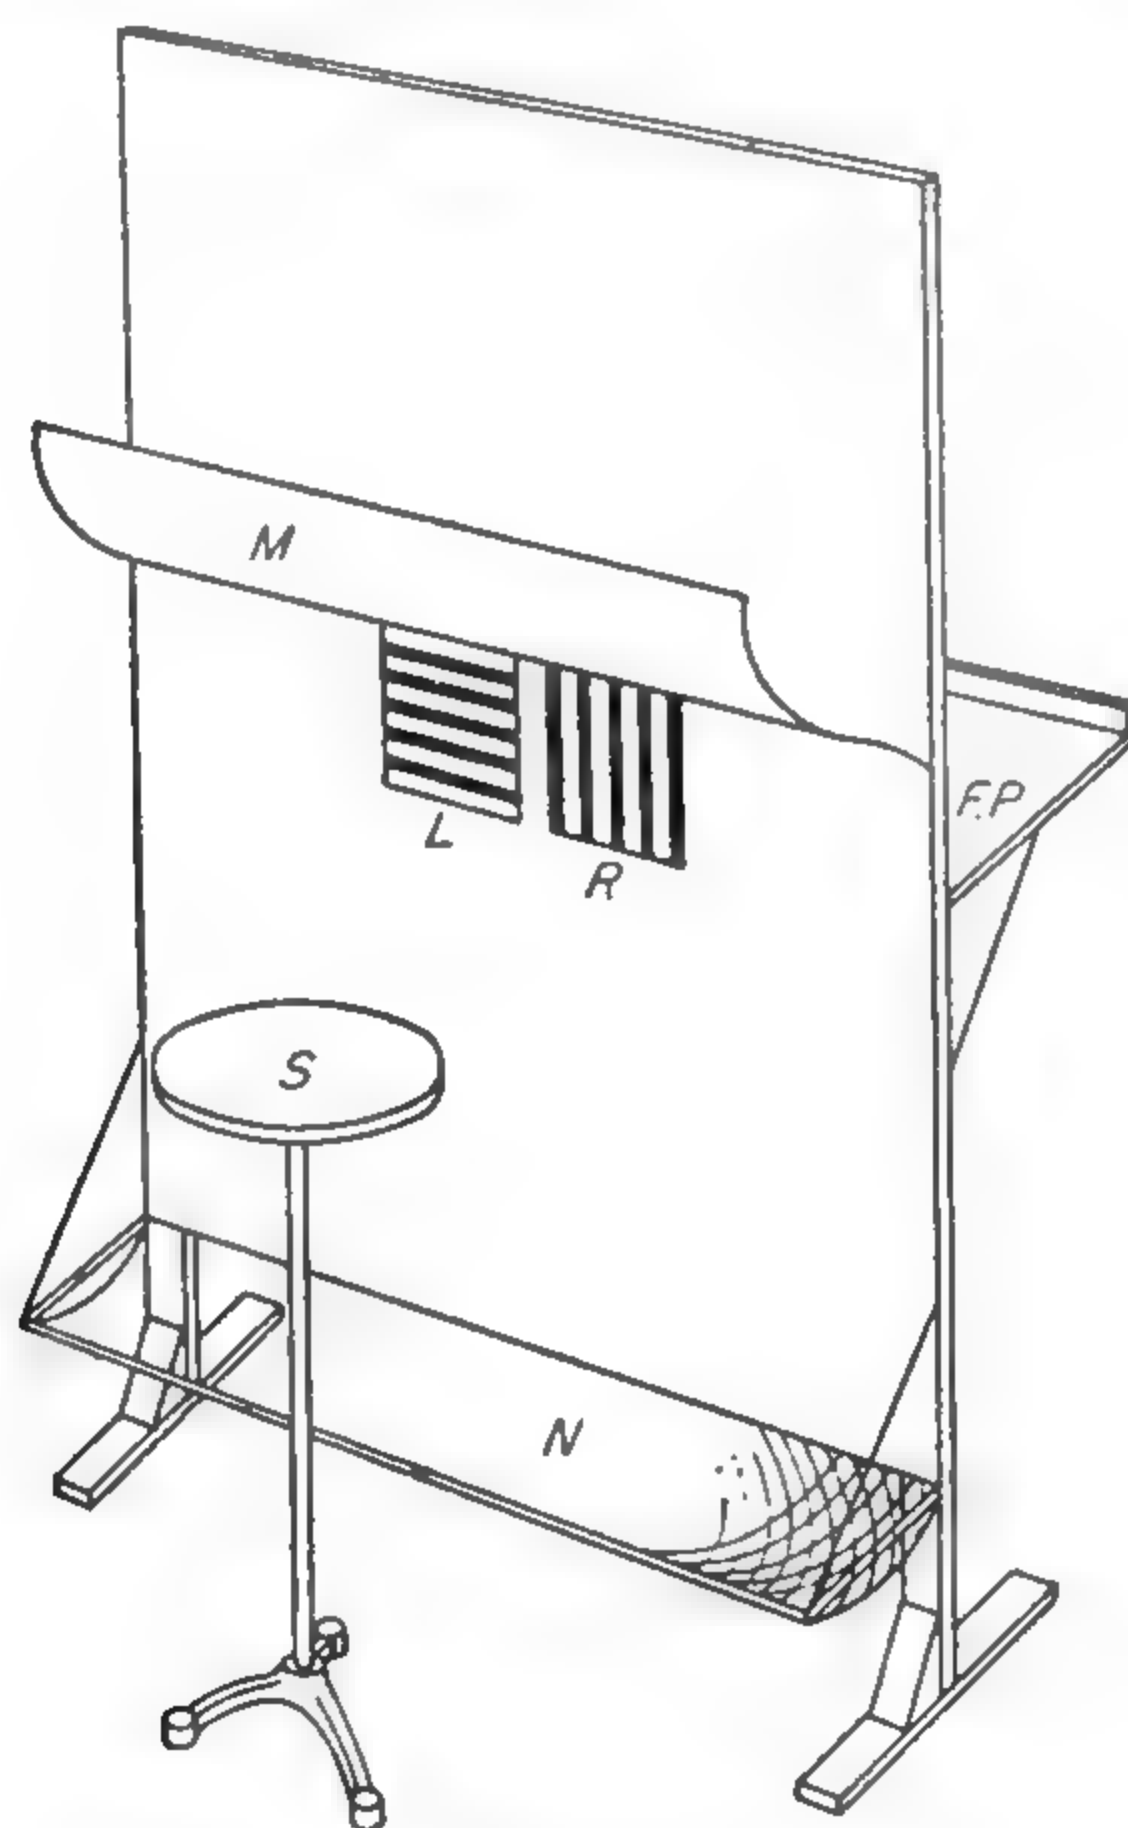


Figure 16. The Lashley jumping stand. The rat jumps from the stand to one door or the other. If the choice is correct, the stimulus card gives way; if the choice is wrong, the rat falls to the net below. (Lashley, 1935.)

runways have been used. In order to study discrimination in animals not ordinarily seen in the psychological laboratory, birds and reptiles, for example, special techniques must be devised. The fundamental rule remains the same however; animals are required to choose between a stimulus that is associated with reinforcement and one that is not (or is associated with some unpleasant stimulus).

The way in which animals learn such discriminations appears simple enough; with increasing frequency the correct alternative is chosen. If the problem is easy the animal learns to make the correct choice in very few trials, and if it is difficult the animal may not discriminate at all. Usually

the difficulty of the problem is determined either by the closeness of the discrimination required or by the complexity of the stimuli presented. Despite the outward simplicity of the learning process, however, the events that go on in the organism must be very complicated. Let us look at some of the attempts to deal with this problem theoretically.

Theories of Discriminative Learning

Most current theories of discriminative learning fall into two classes. One of these places great emphasis upon the gradients of stimulus generalization and upon the gradual accumulation of habit strength to the correct stimulus. The other class stresses the importance of the active problem-solving character of discrimination and the all-or-nothing characteristic of hypotheses that animals might employ in learning to discriminate.

Unfortunately no single theorist of either class has dealt with all of the problems in discrimination learning. Therefore, in stating these theories, we must draw on the work of different theorists. By combining their views in this way, we probably do an injustice to the individual theorist. At the present time, however, there seems to be no more satisfactory way of obtaining a synthesis of these theoretical problems.

The Algebraic Summation and Continuity Theory

The theory of discriminative learning that makes greatest use of the concepts of stimulus generalization and of the gradual accumulation of habit strength is primarily the work of Spence (1936, 1937a, 1940, 1952), though most of his fundamental notions are close to those of Hull (1943, 1952). Let us look at some of the principles these theorists have emphasized.

First of all, these theorists have insisted that gradients of stimulus generalization interact algebraically. Suppose that an animal is trained in a discriminative problem by the reinforcement of response to one stimulus and extinction of response to another. The effects of reinforcement will generalize to all other related stimuli, and so will the effects of extinction. According to this view the strength of the tendency to respond is obtained by subtracting the strength of generalization of extinction from the strength of generalization of reinforcement for any particular stimulus. Thus the following equation expresses the notion:

$$R_a = R - I$$

R_a is the *net* strength of response after the inhibition from extinction is subtracted.

This hypothesis is illustrated in Figure 17. Here, the solid curve represents the generalization of conditioning R ; the dashed curve, the generalization of extinction I ; and the dotted one, the net response strength R_0 after I has been subtracted from R for every stimulus. The exact form of the generalization curve, as well as the values of its constants, is, of course, unknown, but despite this limitation this idea of algebraic interaction leads to important conclusions.

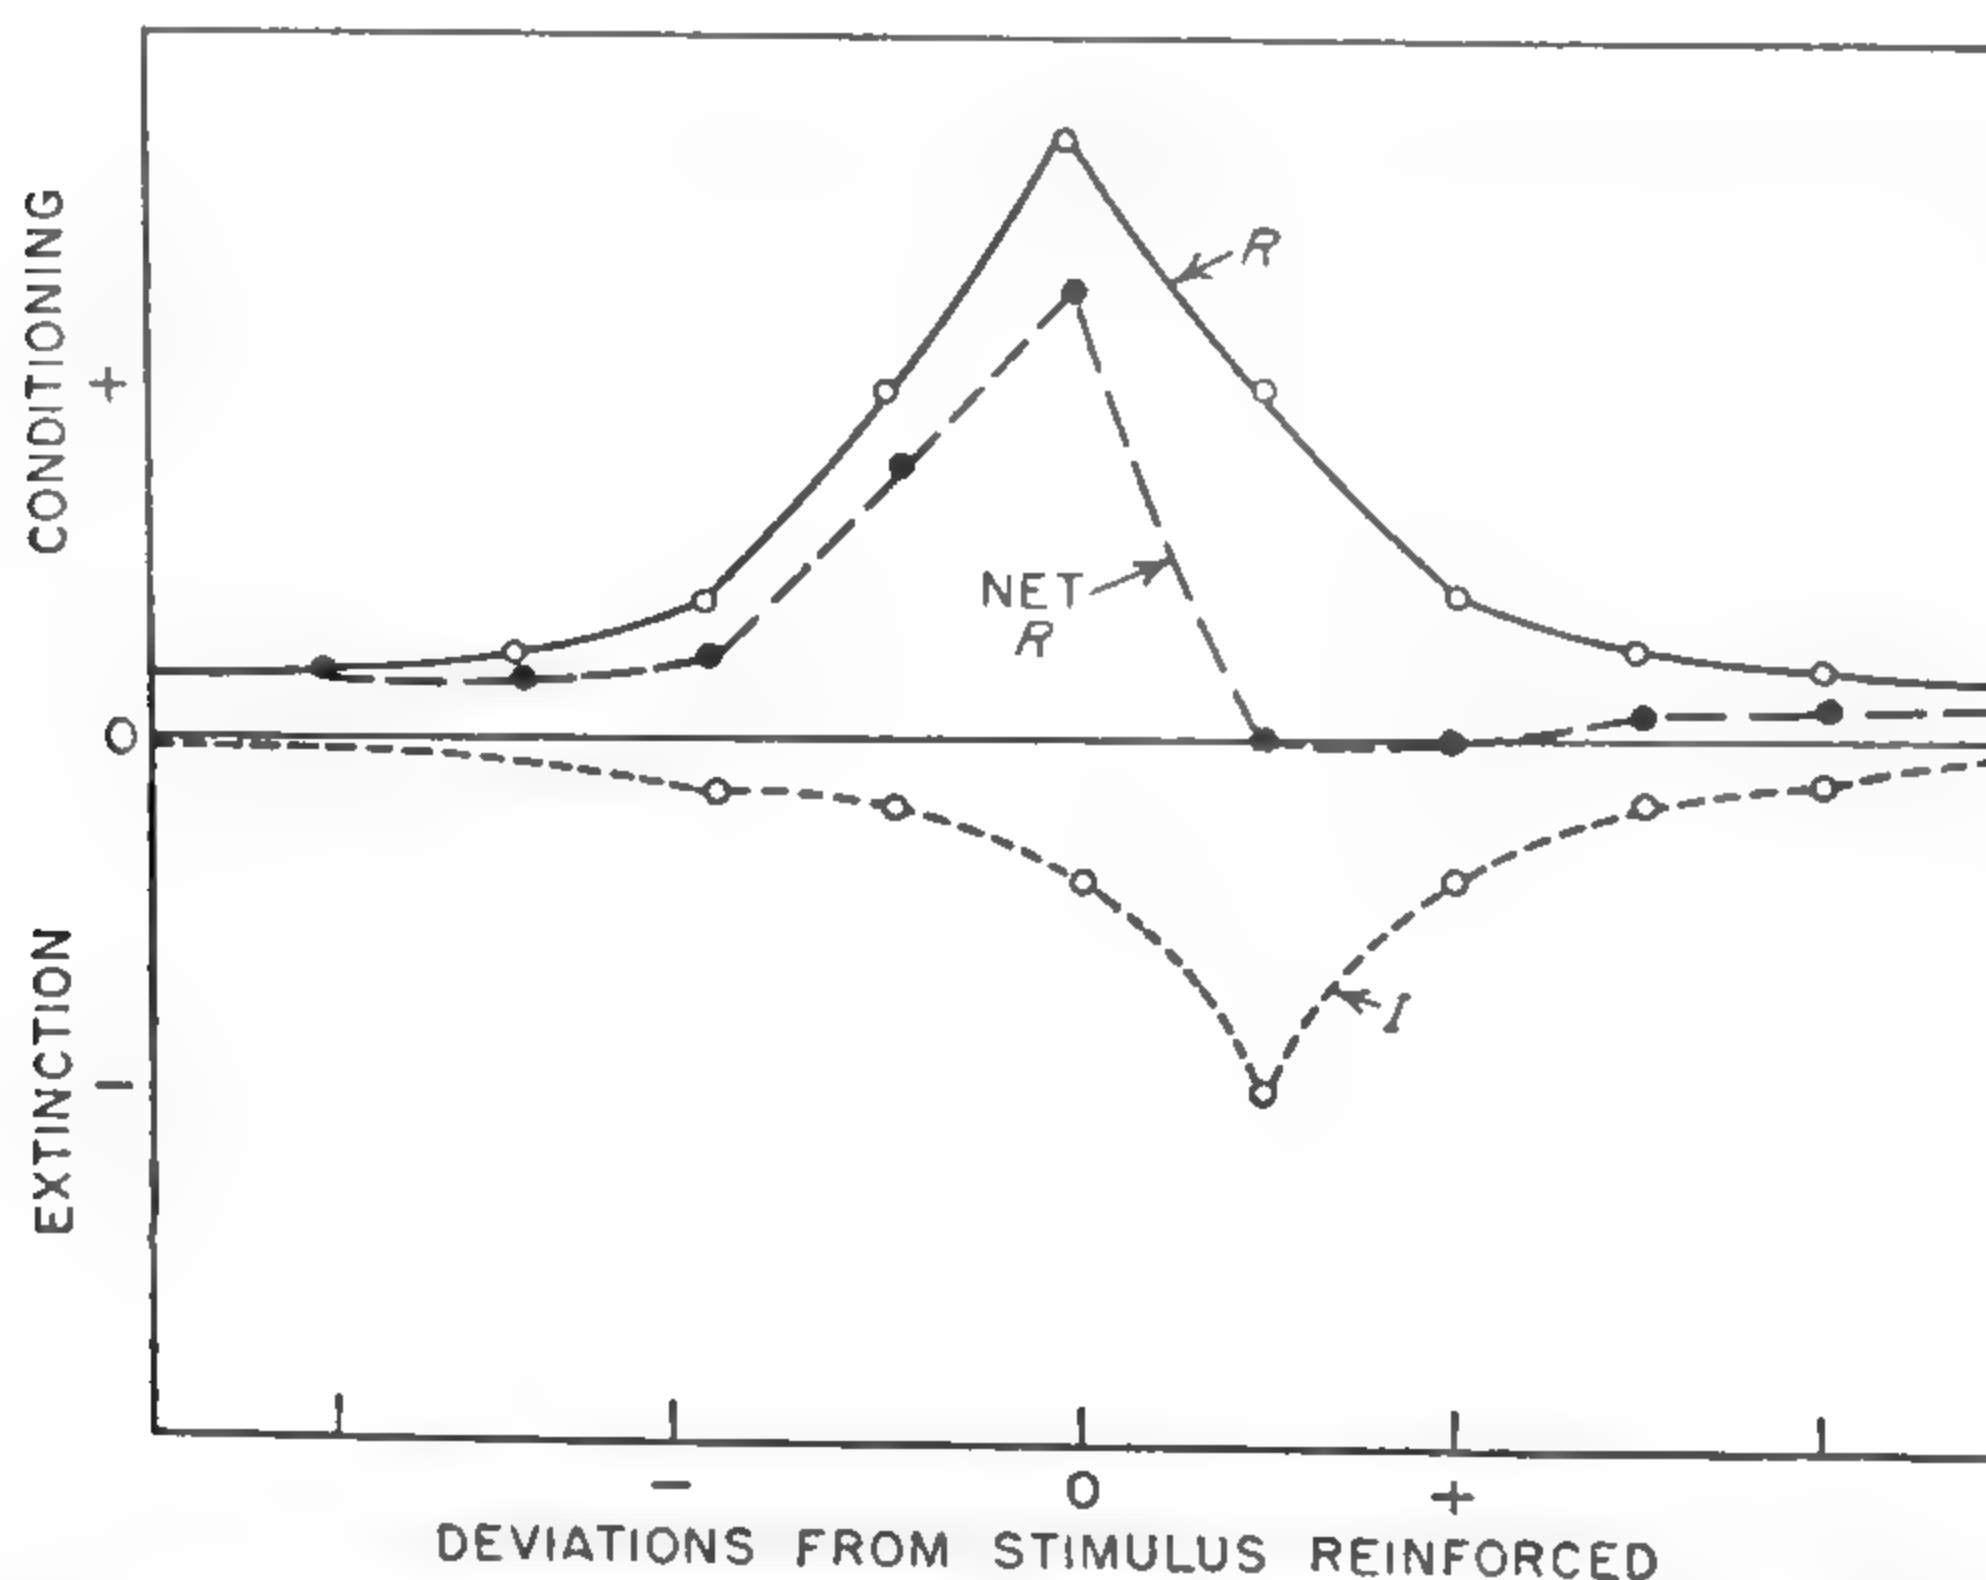


Figure 17. Illustrates the algebraic-summation notion of discrimination learning. Net response strength is the result of the subtraction of the effects of generalized extinction from those of generalized reinforcement. (Suggested from Hull, 1943.)

Spence's analysis of transposition. One of the most interesting by-products of the algebraic summation notion is the way in which it can be used to account for a special phenomenon in discriminative learning known as *transposition*. Transposition is the apparent ability of organisms to make *relational discriminations*. We can best illustrate what these are by example.

Köhler (1915) once trained chickens to respond (with food as reinforcement) to the darker of two gray surfaces. The chickens were never reinforced when they responded to the lighter surface. When this discrimination was well established, the animals were presented with a new choice between the *original* reinforced gray and one *darker* still. The interesting thing was that the animals chose the darker of these two grays, even though they had always been reinforced for choosing the other one. In other words,

argued Köhler, the animals had learned the relationship "darker than" rather than the association of the original stimulus with reinforcement.

This result presents a paradox for the theory of learning, for here is a case in which reinforcement apparently has no effect (or even a negative effect) on choice. It turns out, however, that this example of transposition can be explained by the principle of stimulus generalization. Let us turn to Spence's explanation.

Spence (1937a, 1937b) uses an illustration based on a study of the discrimination of visual size in chimpanzees. He trained animals to discriminate between two squares, 256 square centimeters and 160 square centimeters in size. The response to 256 was reinforced and the response to 160 was extinguished. Then, after the discrimination between stimuli 256 and 160 was established, the values of the stimuli were shifted to 256 and 409. The animals chose 409, a stimulus which had never been reinforced.

The theoretical analysis of this problem is presented in Figure 18. The extinction to stimulus 160 generalized to 256 to some extent; furthermore, the conditioning to stimulus 256 generalized somewhat to 409. To arrive at the comparative value of response strength after discrimination of 256 and 409, we subtract the amount of generalized extinction from the amount of generalized conditioning. As shown in Figure 18, when the net response strength is obtained by this procedure, it is greater for 409 than for 256. Thus, we would predict that, on the average, animals would choose to respond to 409 rather than to 256. This is a prediction of the transposition effect.

The generalization curves in Figure 17 do *not* yield the transposition effect; in this figure the highest net response strength occurs to the stimulus originally reinforced. From a comparison of this figure with Figure 18, we can infer that the prediction of whether or not transposition will occur depends upon the shape, relative location, and slope of the generalization curves. This latitude offered by the selection of arbitrary values (presumably to be determined experimentally) for curves of generalization is, theoretically, both an advantage and a weakness. It is a weakness in that the theory can provide no rational ³ way of exactly stating the quantitative aspects of its assumptions; it is an advantage in that it suggests that transposition will occur only sometimes. The *gestalt* notions from which the transposition problem originally comes imply that transposition will always

³ Spence (1942) does present some empirical constants for a particular form of generalization, but since we have already seen something of the variability of obtained generalization curves, we cannot stretch his curves to fit many cases.

occur. The theory based on generalization gradients suggests that occurrence depends upon how we affect the generalization gradients by training, etc.

The obvious way to *prevent* transposition is to make the positive and negative stimuli very different from each other. A demonstration of what this would do can be seen by moving the curves in Figure 18 apart. In general, experimenters have found that when they move stimuli far apart, there is a failure of transposition (Spence, 1937a; Webb, 1950). Two studies (Kendler, 1950; Ehrenfreund, 1952) have systematically explored the

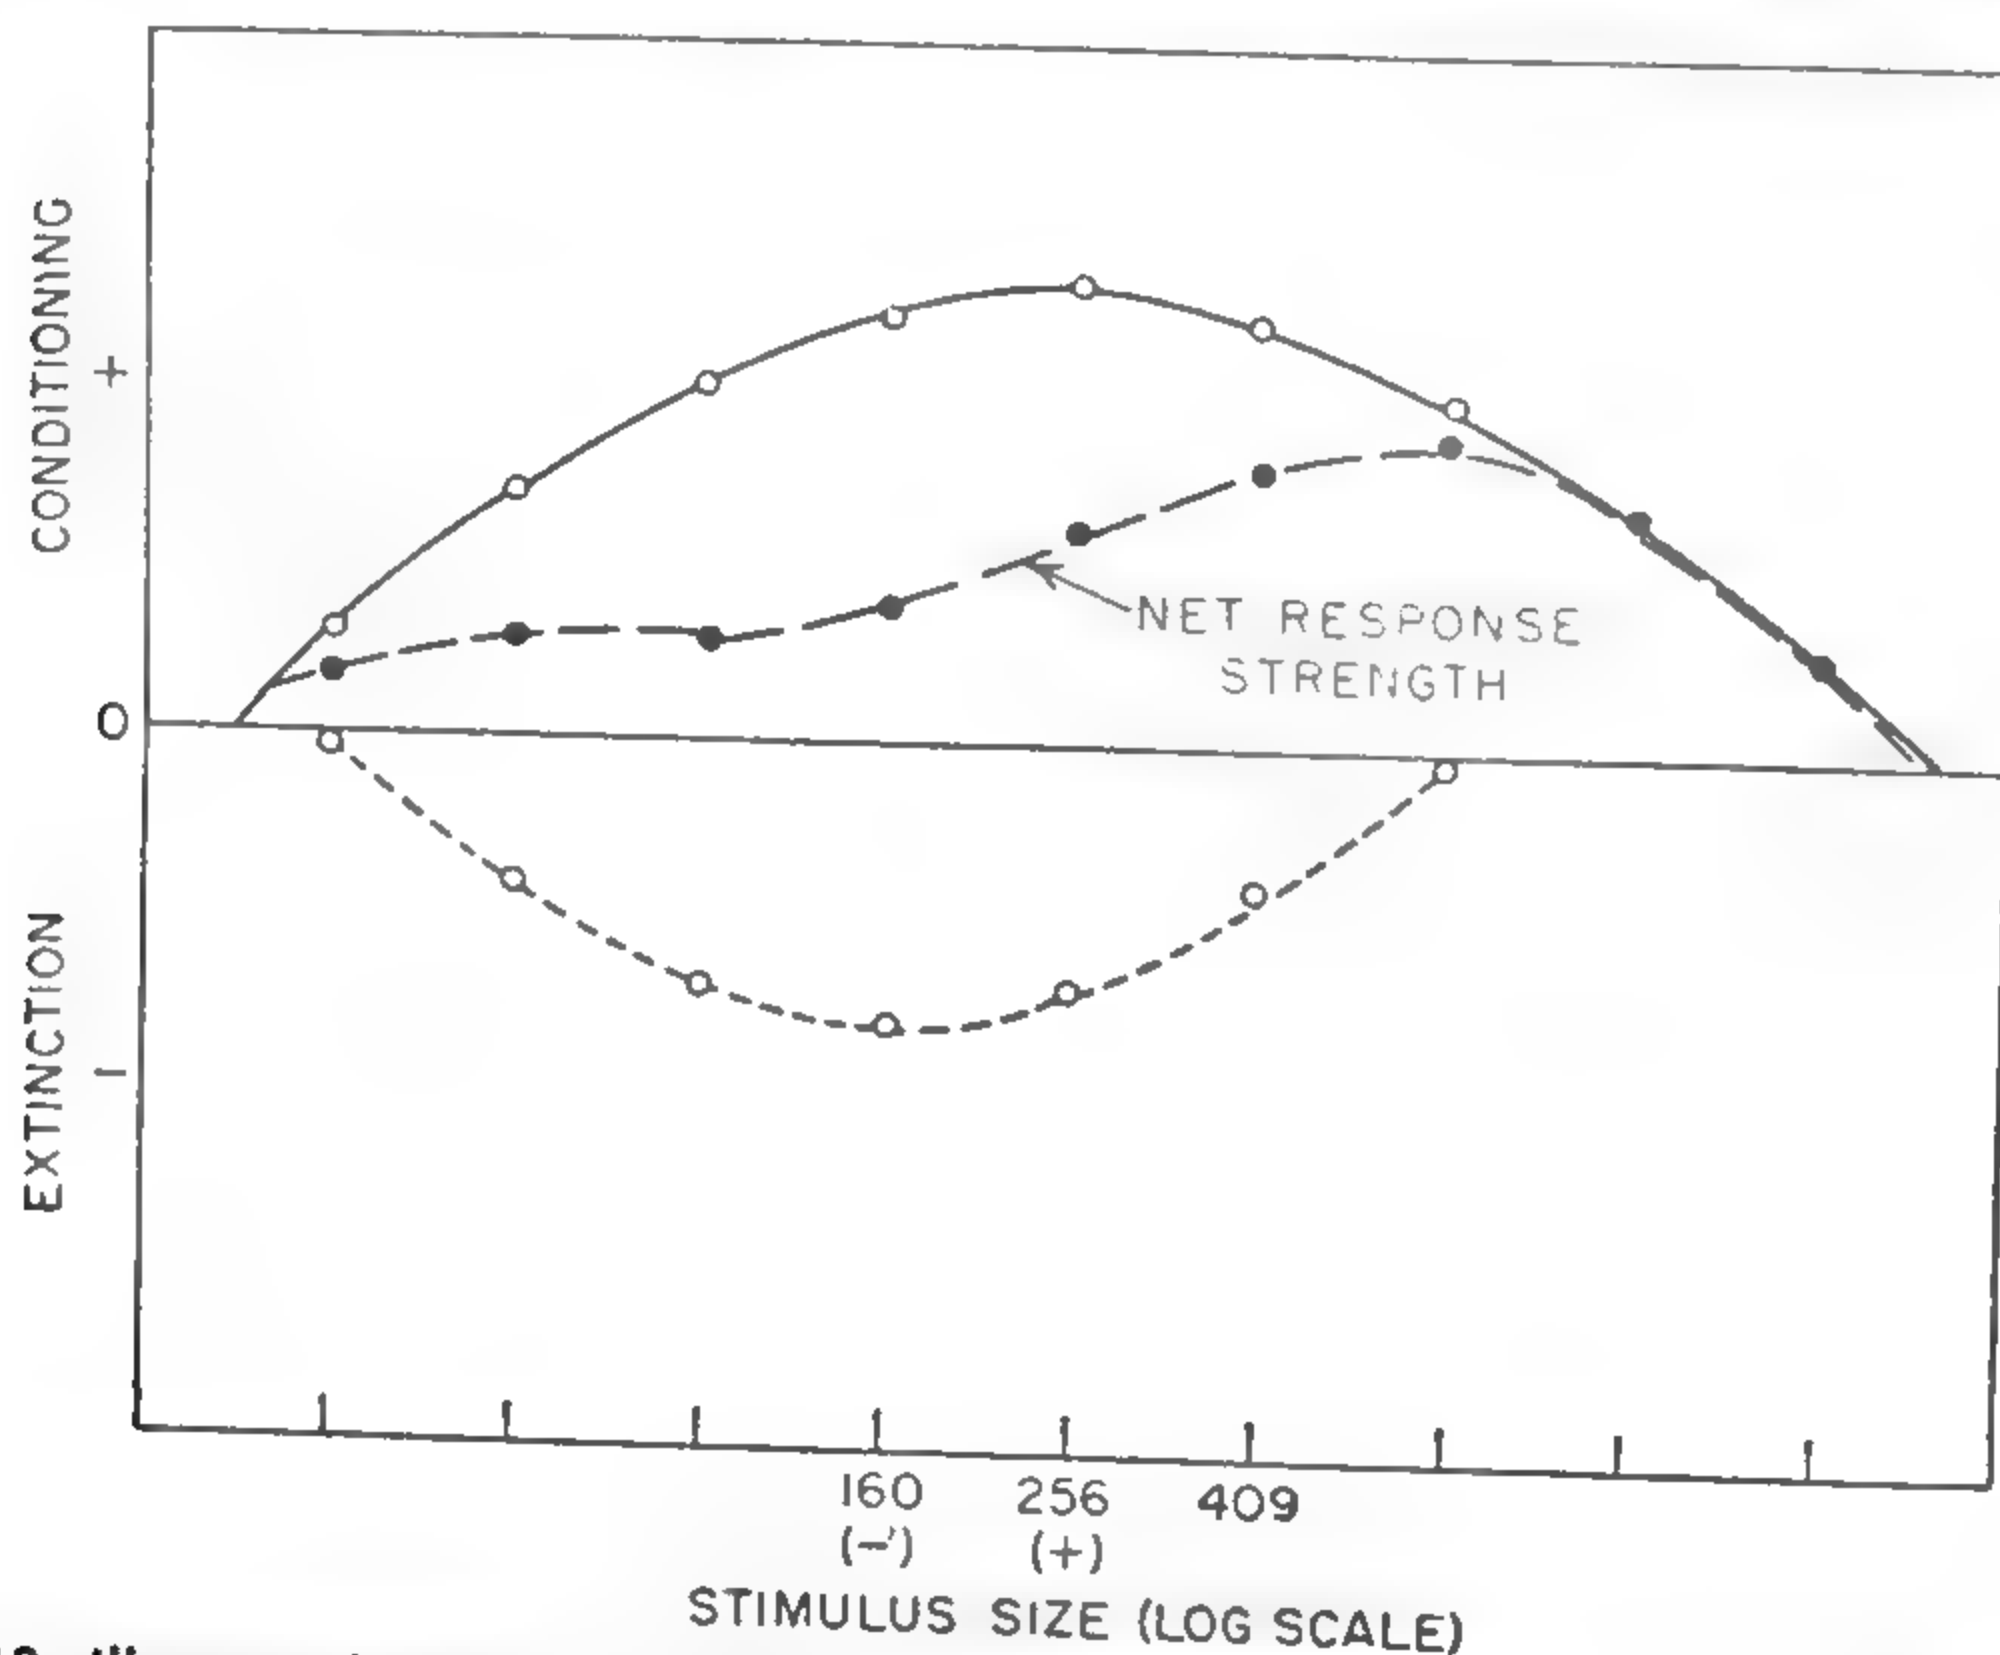


Figure 18. Illustrates how the algebraic-summation notion accounts for an example of "transposition" behavior. The net response strength is greatest, not at the stimulus size specifically reinforced (256), but at larger ones (409, etc.). (After Spence, 1937a.)

effects of increasing the difference between stimuli, and the results are what would be expected from generalization gradients—the percentage of responses showing transposition decreases.

Several experimenters have studied the problem of transposition in children (Kuenne, 1946; Hunter, 1952; Alberts and Ehrenfreund, 1951). Arguments have been advanced to the effect that as children pass from pre-verbal to verbal stages their solution of the transposition problem changes from one based upon generalization gradients to one more like adult formation of "relationship" concepts (Kuenne, 1946; Stevenson, Iscoe, and McConnell, 1955). No critical evidence exists that it is so, but it is an interesting hypothesis, and we shall encounter it later.

Despite the considerable merits of the stimulus-generalization derivation of transposition, we should not lose sight of the fact that sensory experiments have conclusively demonstrated our perception of objects to be determined by the context in which they are embedded (gray compared with black looks different from the same gray compared with white). In this sense, even discriminations based on the subtraction of generalization gradients should be relational; i.e., the generalization gradients themselves will be modified by perceptual relationships between the stimuli. Finally, rats will transpose more readily when the stimuli are presented simultaneously (Baker and Lawrence, 1951). Furthermore, rats will unambiguously make relational choices when they are presented with stimuli that clearly imply relationships. For example, relational choices are more likely if the stimuli are of two brightness levels rather than one (Lawrence and DeRivera, 1954).

Other uses of the algebraic summation principle in generalization. The basic notion used in the prediction of transposition from generalization gradients, namely, the algebraic summation of generalized habit tendencies, has many other applications. It has been generally assumed by those working in the psychology of learning that stimulus generalization is a basic component in the ability of organisms to transfer what they have learned from one situation to another, similar one. The relationships between stimulus generalization and more general problems in the transfer of training involve many difficulties, however, and we must postpone a discussion of some of these. Part of the difficulty, of course, is our inability to specify the exact form of generalization gradients. It should be noted that some progress is being made; Lawrence (1955), for example, has pointed out certain important restrictions that must be incorporated into generalization gradients if they are to account for the existing data on discriminative learning. Lawrence suggests that only certain types of *concave downward* gradients are adequate (see Figure 14).

The Noncontinuity Theory

The theory of discriminative learning we have just looked at places great emphasis upon the gradual accumulation of habit strength and the algebraic summation of gradients of generalization based on reinforcement and extinction. In general this theory has been contrasted with one which emphasizes the problem-solving behavior of organisms in discriminative learning. This is generally called a *noncontinuity* theory, in that it implies that learning a discrimination is not a continuous accumulation of positive and negative habit strength. The essential idea is that animals try out hypotheses

about the discriminative problem they are required to solve. Let us see how this might work.

A rat is trained to jump to one or another window in the Lashley jumping stand. The problem for the rat is to learn to jump to the white window and to avoid the black one. When initially faced with a choice between the white and black, the white window might by chance be on the right side. If the rat jumps successfully to the white window, it might learn either that white is correct or that the right side is correct. Let us suppose that it learns right to be correct. The rat will then work on this hypothesis until it is either punished or extinguished. The rat might then adopt some other hypothesis, based upon a chance contingency of reinforcement and some aspect of the environment. This would continue until the rat happened on the correct hypothesis and kept at it long enough to convince the experimenter that the rat had solved the problem.

In other words, instead of gradually learning a single habit through the accumulation of the effects of reinforcement, the rat would have learned (and unlearned) a whole series of habits analogous to what human beings call hypotheses.

This notion—that animals adopt hypotheses in solving problems—has been advanced by many investigators, most notably Krechevsky (1932) and Lashley (1938, 1942; Lashley and Wade, 1946). The noncontinuity theory has never been entirely expressed by one author, and many more points than the one about hypotheses have been made. Lashley, for example, in nearly all his publications emphasizes the momentary attentional set the animal may have in perceiving a portion of the stimulus field set in front of it by the experimenter.

The theorists who have supported the noncontinuity viewpoint have tended to use experimental problems that arise out of intuition and their own experience. The continuity theorists, on the other hand, have tended to develop experimental hypotheses more rigorously from formal (and sometimes mathematical) principles. In the light of this difference in approach it is not always possible to compare the continuity and noncontinuity theories; but let us look at some of the experiments designed for this purpose and see what the outcome has been.

Comparison of the Continuity and Noncontinuity Theories

The effects of changing the problem. Some of the best evidence on the comparative adequacy of the continuity and noncontinuity theories has come from experiments in which a problem is altered part way through

the training program. This has most frequently taken the form of reversing the discriminative cues. Suppose that we start to train a rat to discriminate between black and white cards. We reinforce the black card and extinguish the white one. Then after a few trials, we reverse the cues, so that the white card is now associated with reinforcement and the black card is extinguished. What effect will this have on the animal's eventual rate of learning the problem? If the animal had been following an erroneous hypothesis (a position hypothesis, say) during the initial phase of the training, we should expect no effect if we reversed the cues, since the animal had not yet begun to associate the black and white cues with the reinforcement and its absence. If, however, each trial adds a small increment both of habit strength to the positive stimulus and of inhibition to the negative stimulus, as the continuity theory would suggest, then reversing the cues ought to lead to negative transfer or interference.

Quite a few experiments which use the technique of reversing cues during the presolution period have been reported. Nearly all of these show that reversing the cues *does* retard learning (for example, McCulloch and Pratt, 1934; Ehrenfreund, 1948; Ritchie, Ebeling, and Roth, 1950; Grice, 1951). If cues are reversed every few trials, animals eventually come to do better on successive reversals. Figure 19 shows that there is interference (reversals result in more errors) the first few times a pair of stimuli is reversed. After several reversals, however, animals come to do better than

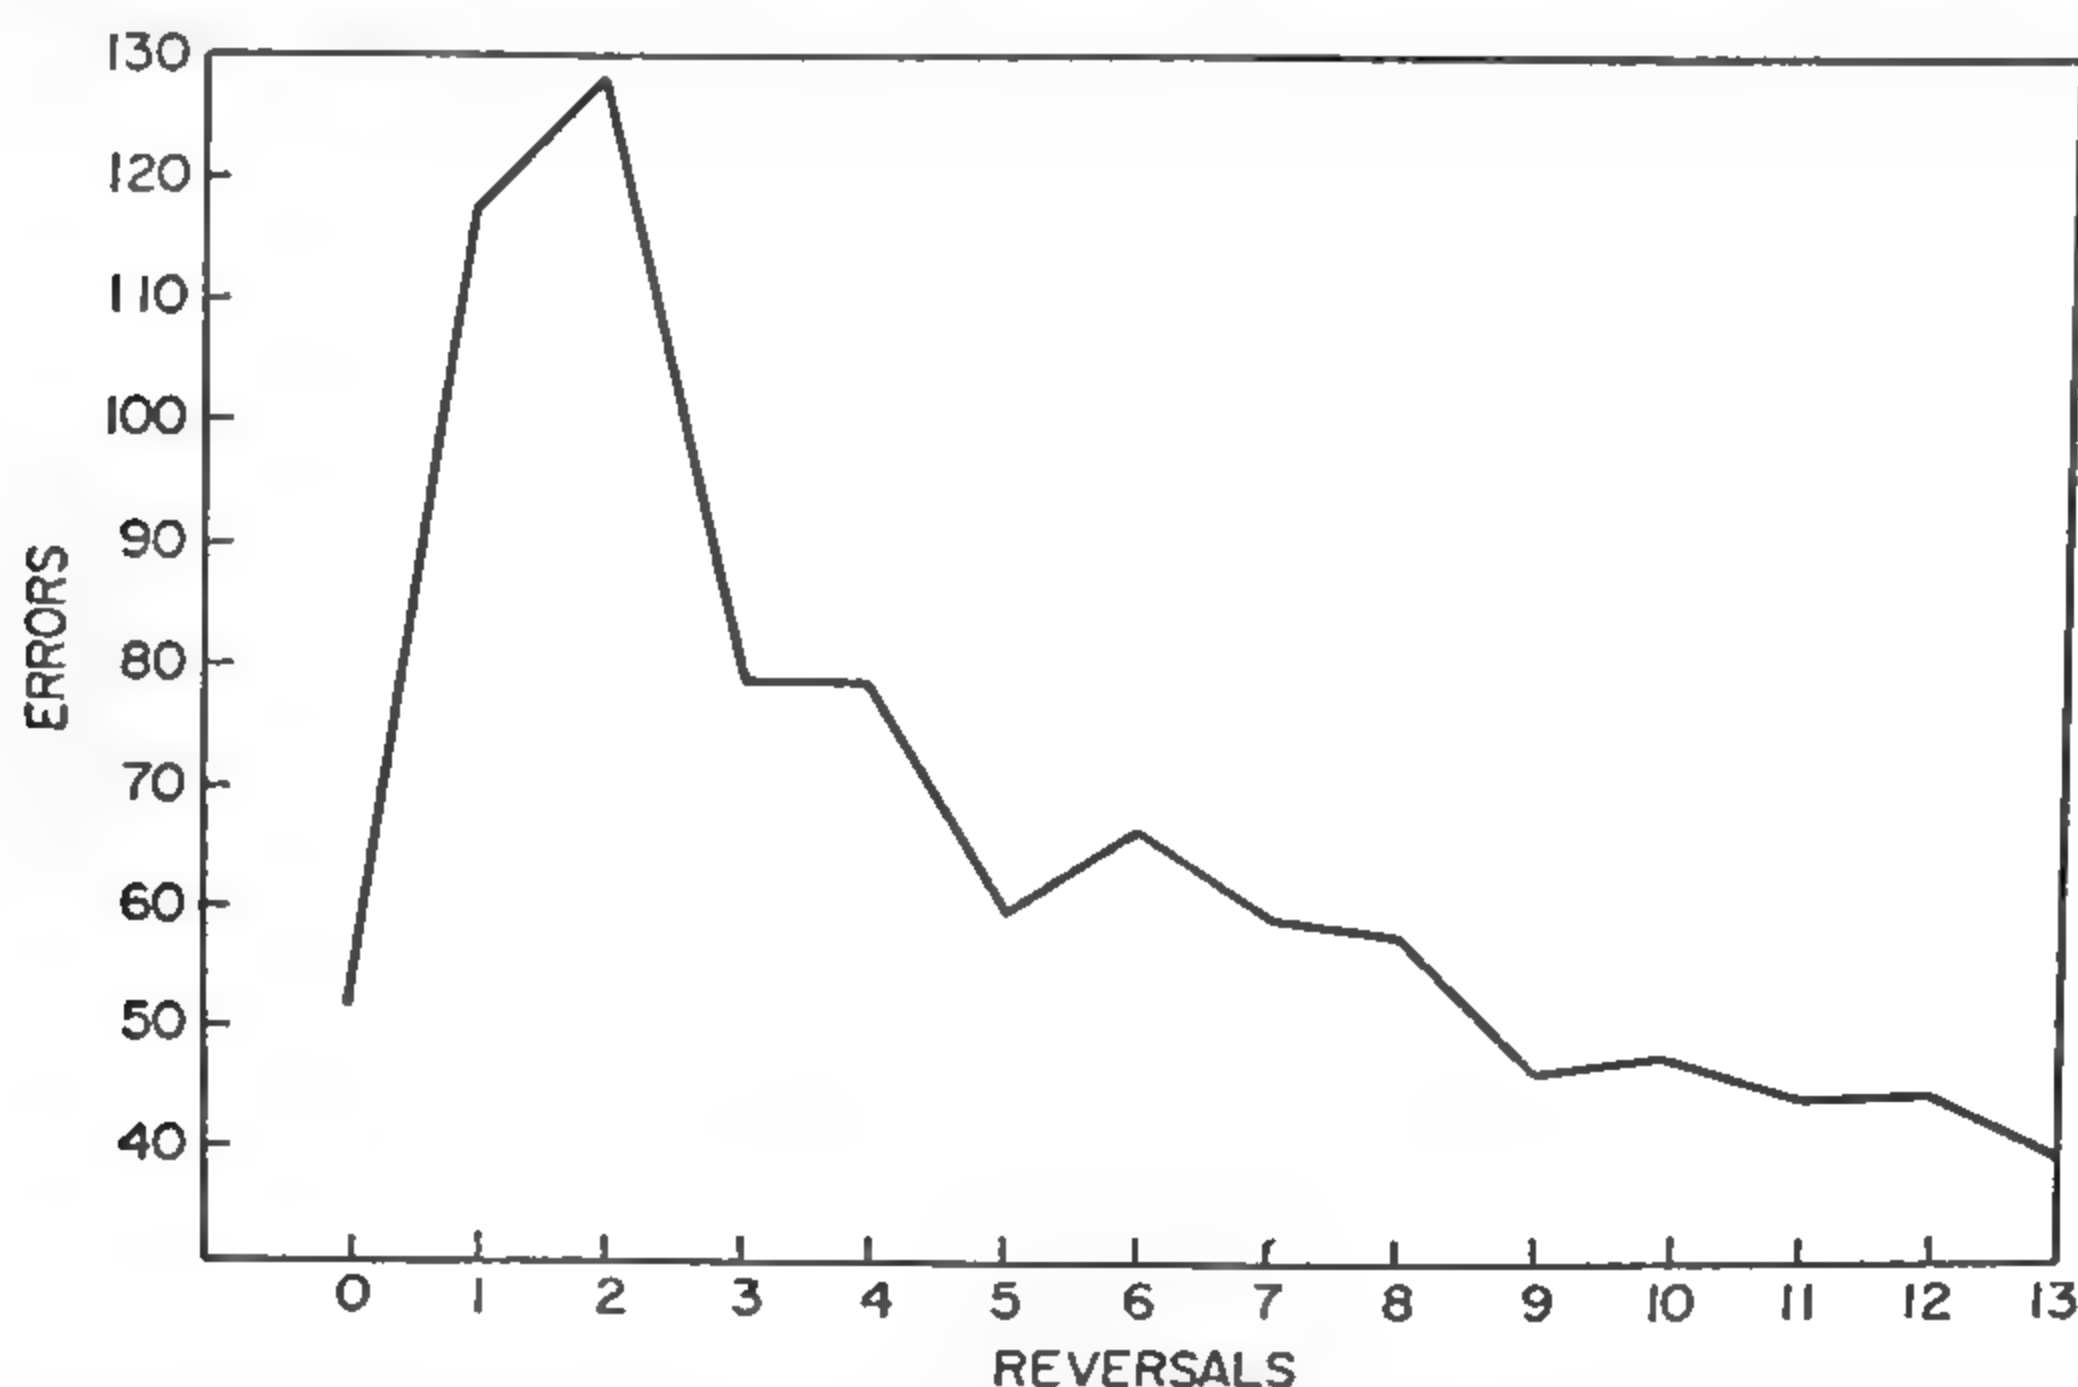


Figure 19. The number of errors made on successive problem reversals. The number of errors on each problem increases rapidly at first, but after a few reversals, errors are less frequent than on the initial problem. (Gatling, 1952.)

they would have done with no reversals at all (Gatling, 1952). This is much like the "learning-set" phenomenon discussed below (also see North, 1950).

It is clear, then, that the results of discriminative experiments in which the cues are reversed favor the continuity view. In addition to these experiments, there are some in which cues are modified rather than reversed during the presolution period and still others in which more than one kind of cue is used. Let us look at some of these experiments.

One that seems to favor the noncontinuity view is an experiment by Lashley (1942). Lashley first trained some rats to discriminate between large and small circles (the large circle was reinforced). Then Lashley continued the training, but this time he substituted a large triangle for the large circle. Thus the animals now discriminated between a large triangle and a small circle. Yet, these animals failed to make the distinction in shape, for when they were tested with the large triangle and the large circle they showed no preference. Furthermore, they showed a preference for a large circle over a small triangle. The results of this experiment suggest that the rats did not learn the incidental cue (shape), but only the cue they had been originally trained on (size).

On the other hand, there are experiments which suggest that rats *do* learn such incidental cues. For example, Bitterman and Coate (1950) showed that if a spatial cue is made incidental to the learning of a brightness discrimination, the spatial cue is learned also. It is obvious that this question is by no means settled. Some experiments that we shall examine later will suggest a way of reconciling the findings that sometimes incidental cues are learned and sometimes they are not.

Successive versus simultaneous discrimination. Another comparison of theoretical importance in discriminative training is that between simultaneously and successively presented cues. (It is not always clear which theory is on which side, however. Compare Grice, 1949, with Spence, 1952.) The problem is especially important to the matter of relational versus absolute discrimination. If animals respond to relationships between stimuli in learning how to discriminate, then successive presentation of stimuli should make the problem more difficult than simultaneous presentation. If, however, animals respond simply to the positive reinforcement of one stimulus and to the extinction of the other, it should not make much difference, all other things being equal, whether the stimuli are presented together or not.

One of the best experimental tests of this problem is by Grice (1949). It is difficult to compare the learning which occurs to simultaneously presented stimuli and that which occurs to successively presented stimuli be-

cause in the latter there are no errors, since the animals have no choice to make. Grice, however, estimated the errors in successively presented stimuli by counting as errors all runs to the positive stimulus with latencies greater than the median latency and all runs to the negative stimulus with latencies less than the median. With this estimate, there was no difference in the learning curves for the two conditions. Thus, Grice concluded, it is not necessary for animals to compare stimuli in order to learn a discrimination.

Spence (1952), however, argues from the absolute theory of discrimination that under certain conditions simultaneous presentation should be easier for animals than successive presentation. He cites data from his laboratory and disputes evidence advanced by Weise and Bitterman (1951) showing simultaneous presentation to be more difficult.

The problem is a complicated one. It is clear, as Nissen (1950) points out in a paper discussing earlier experimental work, that discrimination is to a large extent the establishment of approach tendencies to positive stimuli and of avoidance tendencies to negative stimuli. The experiments on successive versus simultaneous discrimination, however, show that the learning of even a moderately simple discriminative habit is too involved to be easily explained by any existing theory. The experiments by Bitterman and others are complicated, as Spence has remarked, by the existence of patterning between positive and negative stimuli; the extent to which this influences ease of discrimination is unknown. One respect in which Grice's experiment is different from the others is that it is not complicated by stimulus patterning. Thus we should probably accept his experiment as demonstrating that absolute tendencies to approach or avoid stimuli can be established, though relationships between stimuli may be fundamental to real-life animal discrimination.

An evaluation of the two theories. It should be evident by now that experimental data on discriminative learning have far outrun theory. Certainly the absolute, continuity theory is correct in asserting that animals can be trained to have positive responses to certain stimuli without comparing these stimuli to others on the same sensory continuum. In the process of training, even though animals give no evidence of having grasped the discrimination, experience at the task has effects which can be detected by negative-transfer tests. On the other hand, there is also abundant evidence that animals do not simply react to any stimuli that come within range; they are selective in what they attend to. Furthermore there is good evidence that animals test preliminary "hypotheses" about the nature of the

problem.⁴ This, of course, is reflected in the experimental literature; this is why there are no unequivocal data to support the extant theories. From the vantage point we have at the moment, it appears as though an adequate theory will be a synthesis of many current notions. The idea of stimulus generalization, or something like it, is important and will probably be retained as an essential element in theory making. In addition there will probably be notions describing more adequately the selective nature of attention and hypothesis in discriminative learning.

THE DEVELOPMENT OF DISCRIMINATIVE CAPACITY

An important problem that has received much experimental attention is that of the development of the capacity for discrimination. This has been studied in two ways. In the first place, there has been much experimental work on the acquisition of distinctiveness of cues in discrimination, and secondly, there has been an elaborate experimental program to study the gradual improvement in discrimination that comes from learning how to discriminate. Let us first consider the problem of acquiring distinctiveness of cues.

The acquired distinctiveness of cues. Much of the theoretical analysis of discriminative learning has incorporated the assumption that cues are always learned to particular responses—that learning is a simple matter of stimulus-response bonds. Nearly all theorists agree, however, that at best this is an oversimplification, and recently there have been many experiments designed to show that discriminative learning involves some mediated or central associations. Some of the best of these experiments have been on the acquired distinctiveness of cues.

Lawrence and his associates (Lawrence, 1949, 1950; Bauer and Lawrence, 1953) have established that when animals learn to discriminate between cues, they can transfer this discrimination to new situations in which entirely different instrumental responses are required. Thus, the distinction between cues is not correlated exclusively with the responses used in discrimination learning, and therefore discrimination learning cannot be simply a matter of stimulus-response bonds. Thus, in addition to consider-

⁴ Much of the evidence for this comes from a phenomenon called "Vicarious Trial and Error." This refers to the fact that animals engage in vacillation and indecision at choice points in discriminative learning. Woodworth and Schlosberg (1954) have presented a good review of the effect.

ing the factors of positive and negative stimulus generalization in any complicated learning problem, we must also consider the extent to which particular cues acquire distinctiveness for animals (see Bauer and Lawrence, 1953):

An important outcome of this work has been to point toward the solution of some of the issues in the continuity-noncontinuity controversy. Lawrence and his associates (Lawrence and Mason, 1955; Goodwin and Lawrence, 1955) have suggested that in complicated discrimination problems, animals can learn several habits simultaneously. First of all they can learn to "orient" toward a particular kind of stimulus (black versus white cards, etc.), and secondly they can learn the actual discrimination (white is reinforcing). Suppose, however, there are other cues present, such as the height of a hurdle over which the animal must jump. If the experimenter changes the problem so that the height of the hurdle is the cue, the animal need not necessarily unlearn the black-white discrimination. All that is necessary is that the *orientation* towards brightness be extinguished. If this is so, the animal can learn the new problem without disturbing the old discrimination. Thus, in complicated discrimination situations it is possible to obtain behavior in which the rat appears to engage in particular hypotheses without disturbing other hypotheses. Notice that this derivation of a complex effect in discrimination learning is built upon the simple mechanisms of reinforcement theory. Lawrence has devised many ingenious experiments which point to the way in which problem solving and relational discrimination can be derived from simple mechanisms. The most important step is the abandonment of the simple stimulus-response paradigm. Once we grant the possibility of connections between stimuli and nonovert mediated processes in the animal, it becomes possible to derive complicated effects in discrimination learning from simple premises.

Reid (1953) has presented the same notion, namely, that in addition to learning a discrimination animals also learn something about the nature of discriminative learning, and this can account for the development of non-continuity behavior out of continuity-like processes. Reid demonstrated the adequacy of this idea by showing that there was more rapid learning of a discrimination reversal after a long period of overlearning on the original problem. Reid points out that Spence and others have suggested the need for a *receptor-orienting* act, or something like it. In simpler discriminative experiments receptor orientation, or learning to notice particular cues, is unimportant because only a limited number of cues are presented to the animal. When conditions become more complicated, however, this orienting

becomes a more important determinant of behavior. It is significant that the experimental evidence shows this orienting is learned in much the same way that other kinds of discrimination are learned.

Discrimination-learning sets. The examples of discrimination discussed thus far have been those in which animals obviously have had some difficulty in learning. Yet it is possible to find examples of discrimination in animals in which there is practically no learning, because the animal solves the problems immediately. In some experiments, for example, two stimuli are presented to monkeys; one of the stimuli, say, is a red cylinder and the other is a blue pyramid. Under one of these objects is hidden a peanut. When these stimuli are exposed, the monkey carefully examines each and chooses to look under one. If, by chance, it looks under the wrong one, on the very next trial it looks under the other one. The discrimination is perfect, and it is only a matter of chance whether the monkey discovers the peanut on the first or on the second trial.

Immediate and perfect solution of this and much more difficult problems comes from a history of learning to make discriminations of this sort. As animals solve successive discriminative problems, they show an orderly and gradual improvement in this ability. Such an improvement in ability Harlow has called the formation of *learning sets* (Harlow, 1949). The demonstration of this gradual improvement is extremely important because it suggests a direct continuity between the most primitive examples of discriminative learning and the insightful behavior of sophisticated animals. One kind of behavior arises out of the other.

An analysis of the way in which this happens has just begun. One important point has been made (Riopelle, 1953)—transfer from problem to problem based upon stimulus generalization of cues associated with successful solution becomes less and less as additional problems are solved. In other words, animals learn that they cannot rely on stimulus generalization from one problem to the next to lead them to rapid solution of successive problems. This means that stimulus generalization of the sort we discussed earlier either becomes completely suppressed or radically altered in character. Warren (1954) points out that test-wise monkeys approach each recombination of stimuli as a *new* problem however many times the specific stimuli required in the discrimination had been differentially rewarded. Thus, even if one of the stimuli were much more frequently rewarded than the other in the past, the animal would not tend to choose this stimulus more often. Warren suggests that this means that the analysis of discriminative learning according to the excitatory and inhibitory strengths built up

to particular stimuli as the result of reinforcement and extinction does not apply to animals that have learned many problems.

Thus as a result of long experience at discrimination, animals learn something new. They are no longer limited to the information that red objects or square objects are associated with reinforcement; they learn that one of two stimulus objects will be reinforcing or that one that was not reinforcing on the previous trial will be reinforcing on the next trial. Such learning is truly the solution of a problem, and we shall reexamine discriminative learning sets when we consider the role of learning in problem solving.

PERCEPTUAL LEARNING

The influence of learning on perception. Almost as old as philosophical inquiry itself is the question, Is man's ability to perceive objects of the world innate or is it the result of learning? In modern psychology this question had long been dormant because nobody seemed to be able to suggest a reasonable or even testable solution to the question. In recent years, largely because it was raised in a new context (Hebb, 1949), there has been a revival of interest in the matter. Specifically, the question has been raised as to whether the normal visual and tactual perception of shape and pattern depends upon learning or whether it is an inherent property of the nervous system.

In the first experiments aimed at exploring this question, apes were reared in darkness from birth to the time of testing. These apes, deprived of all visual experience, were tested on visual discrimination problems. The results showed an astounding deficiency in visual form perception; these findings were not conclusive, however, because there was evidence of a deterioration in the physical capacity of the eyes because of the long period of disuse (Riesen, 1950).

Several subsequent experiments have confirmed the earlier observations, and in these the possibility of deterioration of the sense organs was fairly well ruled out. In one (Siegal, 1953), doves were reared with translucent head covers. These covers permitted light to come through (thus not depriving the retina of stimulation) but did not allow definition of visual form. Birds reared with these covers took much longer to learn simple form discriminations than did normal birds. It seems certain that experience was a normal prerequisite for visual perception in these birds, though there is some question as to whether a lack of visual-motor coordination was not the basic factor.

In another experiment (Nissen, Chow, and Semmes, 1951) a chimpanzee was reared deprived of tactual, kinesthetic, and manipulative experience. The results were much the same as those found in the experiments with the ring doves. This animal was grossly inferior to normal ones in the ability to learn tactual discriminations. Again it is not clear whether the deprivation in experience resulted in a purely sensory deficiency or whether the deficiency was one in sensory-motor coordination.

These experiments point to the great importance of experience in the performance of even the most simple perceptual-motor acts. The deficiencies that result from deprivation are not the result of failure to learn how to discriminate generally, since the animal is free to perform discriminations in its other senses. The deficiency is specific to the sense deprived. If it turns out, as Hebb believes, that the deficiency is a result of the failure of the nervous system to *learn* the nature of forms in the perceptual world, then the role of infantile learning is enormously more important than anyone has ever thought. There is strong incidental evidence to support Hebb's views, since it is difficult to understand how the perceptual constancies could arise unless perceptual relations are learned.

Response versus Place Learning

Experiments on infantile sensory deprivation and perception suggest that learning of perceptual relationships can occur without the intervention of motor acts. Other sources supply good evidence for this. As we have seen, some of the experiments on complex discriminative learning suggest that it is necessary to assume that some discriminations are learned without being attached directly to motor responses. In the preceding chapter it was also pointed out that instrumental conditioned responses could be extinguished without the performance of the originally learned response. Another problem arising from this question of pure perceptual learning concerns the nature of what animals, particularly rats, learn in mazes. A pure stimulus-response theory would say that rats learn chains of discriminations attached to chains of responses. Thus, at a particular signal, the rat learns to turn right. According to this view, rats would never really learn locations, only responses.

A considerable literature exists, however, purporting to show that rats do learn more than simple chains of responses; they do indeed appear to learn locations. In one experiment (Tolman, Ritchie, and Kalish, 1946), rats were trained to run a rather peculiar maze. After the rats had learned the pattern of the maze, new pathways were supplied and the old, direct

route to the goal was blocked. Most of the animals chose new pathways, which suggested they had learned the general spatial location of the goal. In another experiment, Galanter and Shaw (1954) showed that place learning is demonstrated in rats if there are many environmental cues, while an earlier study (Thompson and Thompson, 1949) suggested that place learning is absent if spatial cues are minimized.

It is clear from these and other studies that rats can learn "places" as well as chains of stimulus-response associations. It is probable that the highly organized perceptual learning implied in the learning of "place" relies on the gradual development of more primitive discriminative learning. It probably depends upon the growth of perception implied in the experiments on early sensory deprivation. It also probably hinges on the development of "learning sets" out of multiple experiences with simpler stimulus-response discriminative problems. Thus the learning of "place" would appear to be a higher-order example of perceptual learning.

Tolman (1932, 1951) has long insisted upon the importance of perceptual learning in the behavior of animals and men. Despite Tolman's belief, there has been relatively little analytical work directed toward finding out about the characteristics of this kind of learning. Perhaps an extension of the work on acquired distinctiveness of cues and on learning sets will help us to understand perceptual learning more adequately.

SOME CONCLUDING REMARKS

Discriminative learning is an important and fascinating business, and there are many problems we have not discussed. For one thing we have not said much about discriminative learning in adult human beings. Many of the experiments on animals have been duplicated at the human level. For example Loess and Duncan (1952) showed that successive presentation of stimuli generated poorer learning than simultaneous presentation, if the discrimination was difficult. Another investigator (Walk, 1952) showed that reversal of a discriminative problem resulted in a decrement in human just as in animal subjects. Despite the fact that the outcomes of these experiments are similar to those conducted with animals, we suspect that discrimination in human beings is much modified by verbal and symbolic processes. In a later chapter on problem solving we shall return to some of these questions concerning human discriminative learning.

More important is the demonstration of gradients of stimulus generalization in voluntary human behavior. Stimulus generalization is basic to the

ability of organisms to transfer what they have learned from one situation to another. The relationships between stimulus generalization and even simple transfer are quite complicated, however, and we must postpone a discussion of these relationships until other important factors in human learning have been examined. Suffice it to say at present that stimulus generalization of verbal responses to symbols and stimuli has been well established in human learning (Yum, 1931; Gibson, 1941) and in motor behavior (Brown, Bilo-deau, and Baron, 1951). These demonstrations are important, for they enable us to make confident use of the notion of stimulus generalization in human learning.

Unfortunately, there are many other questions about discriminative learning that we shall have to ignore. One of the most interesting concerns the relation of capacity for discriminative learning to the natural behavior of animals and their mechanisms of survival. If this were a book about animal behavior rather than a book about the processes of learning, we should study this question in detail. Let us, for now, take an example of the importance of discrimination in natural animal life as an indication of the role of these factors in survival.

There is a variety of caterpillar that survives because of the keen discriminative learning of its natural enemy, birds. This caterpillar, the larval stage of the Cinnabar Moth, is brilliantly striped in black and yellow. Every young bird tries one of the caterpillars, and rejects it because, apparently, it is highly distasteful (Windecker, cited in Tinbergen, 1953). Thereafter the birds leave this caterpillar alone. Thus, the caterpillar sacrifices some of its kind in order to instruct birds in a discrimination highly useful to itself. Certain species of insects, moreover, mimic the appearance of distasteful insects, and therefore rely, to the extent they resemble the creatures they mimic, on stimulus generalization (Tinbergen, 1953).

As you can see, there is a vast area in natural history in which the notions derived from laboratory experiments on discriminative learning might well be put to use. There is, indeed, a recent movement to put to use in wildlife preservation and control some of the principles obtained from the laboratory study of animal behavior (Smith and Geis, 1956).

CHAPTER 5

MOTIVATION AND LEARNING

For a long time psychologists have made a distinction between the things we learn, our habits, and the things which prompt us to use these habits, our motives. Motives provide the internal impetus behind behavior (needs or drives) and the direction the behavior takes (goals); learning provides the means by which the organism finds a way to reduce or eliminate need by coming into contact with an appropriate goal.

Motivation is a construct, not something directly observable in animal behavior. In other words it is an idea we invent to describe the impetus that makes animals active. Presumably it would be possible to invent a satisfactory theory of behavior that never mentioned the word *motive* or another word like it, but no such theory exists at the present time, and, indeed, psychologists have used motivational constructs very freely. Nearly all psychologists agree that it is necessary to make use of a construct that emphasizes the energy behind behavior. In this chapter we shall discuss the relationships between the motivational constructs and learning. In doing so, we shall view motivation as the energizer of behavior and learning as its guide.

MOTIVATION AND PERFORMANCE

Motivation has at least two fundamental components, the need state and the goal, or external stimulus, that tends to eliminate or reduce that need. That is to say, for nearly every need there is some goal that will satisfy or reduce the need. Needs may be cyclic, or periodic, in character, in which case they usually depend upon changes in the internal, physiological balance of the organism. Hunger and thirst are such cyclic needs. Not all needs are

governed by internal changes, however. Some exist at a more or less constant level, awaiting only the occurrence of the appropriate external stimulus to set off the appropriate behavior. The avoidance of pain is such a need, and in all animals there are many other such needs, needs we ordinarily characterize as instincts.

When a need is aroused, either by internal changes or an appropriate external stimulus, the result is usually that the animal performs some behavior designed to reduce or eliminate the need. Thus, when an animal finds its leg in the jaws of a trap, it struggles to free itself; or when it is thirsty, it finds its way to a familiar watering place. A good many, though not all, of these patterns of behavior are learned. Furthermore, animals learn to anticipate these need states before they actually occur, and to act accordingly. Thus, animals learn to *avoid* traps, and they sometimes find their way to food or water when they are not hungry or thirsty, or at any rate they seem to learn to stay in the vicinity of water and food supplies.

Activity and motivation. The classical theory of needs says that they arouse animals to activity. As a matter of fact, there are some well-established relationships between needs such as general hunger and activity. Activity increases directly as a function of the number of hours that rats have been deprived of food (Siegal and Steinberg, 1949). There is some point, when the deprivation becomes severe enough, at which activity drops off (Skinner, 1938). Rats, when allowed to exist upon a free feeding schedule, will show regular cycles of activity and eating. A period of quiescence is followed by some activity which culminates in eating (Richter, 1927). Eating itself is followed by a period of subnormal activity (Finger, 1949).

It has generally been assumed that all of these effects associated with the arousal of need for food are direct, unlearned effects of the physiological consequences of deprivation. Recently, however, several investigators have suggested that the increased activity and other things associated with hunger are partly the result of learning. The cycle of quiescence-activity-eating-quiescence, for example, does not occur in rats when their environment is carefully held constant (Baker, 1954). Furthermore, activity can be induced by sudden changes in environment (Campbell and Sheffield, 1953), and there is strong evidence that such activity is, in part, the result of a discrimination based upon the cues associated with eating (Sheffield and Campbell, 1954). Some of the ordinary activity-arousal effects of such needs as hunger are partially the result of learning and experience. This is one demonstration of the fact that the behavior patterns aroused by needs are modified by learning.

It is well known that other conditions of need, such as sex, maternity, lowered temperatures, have specific and general influences upon activity. That is to say, there are many different organic conditions that produce changes in the general activity level and that induce animals to perform both unlearned and learned acts specific to each organic condition. Most of these specific effects are not of critical importance to the psychology of learning, however, and we shall largely ignore them.

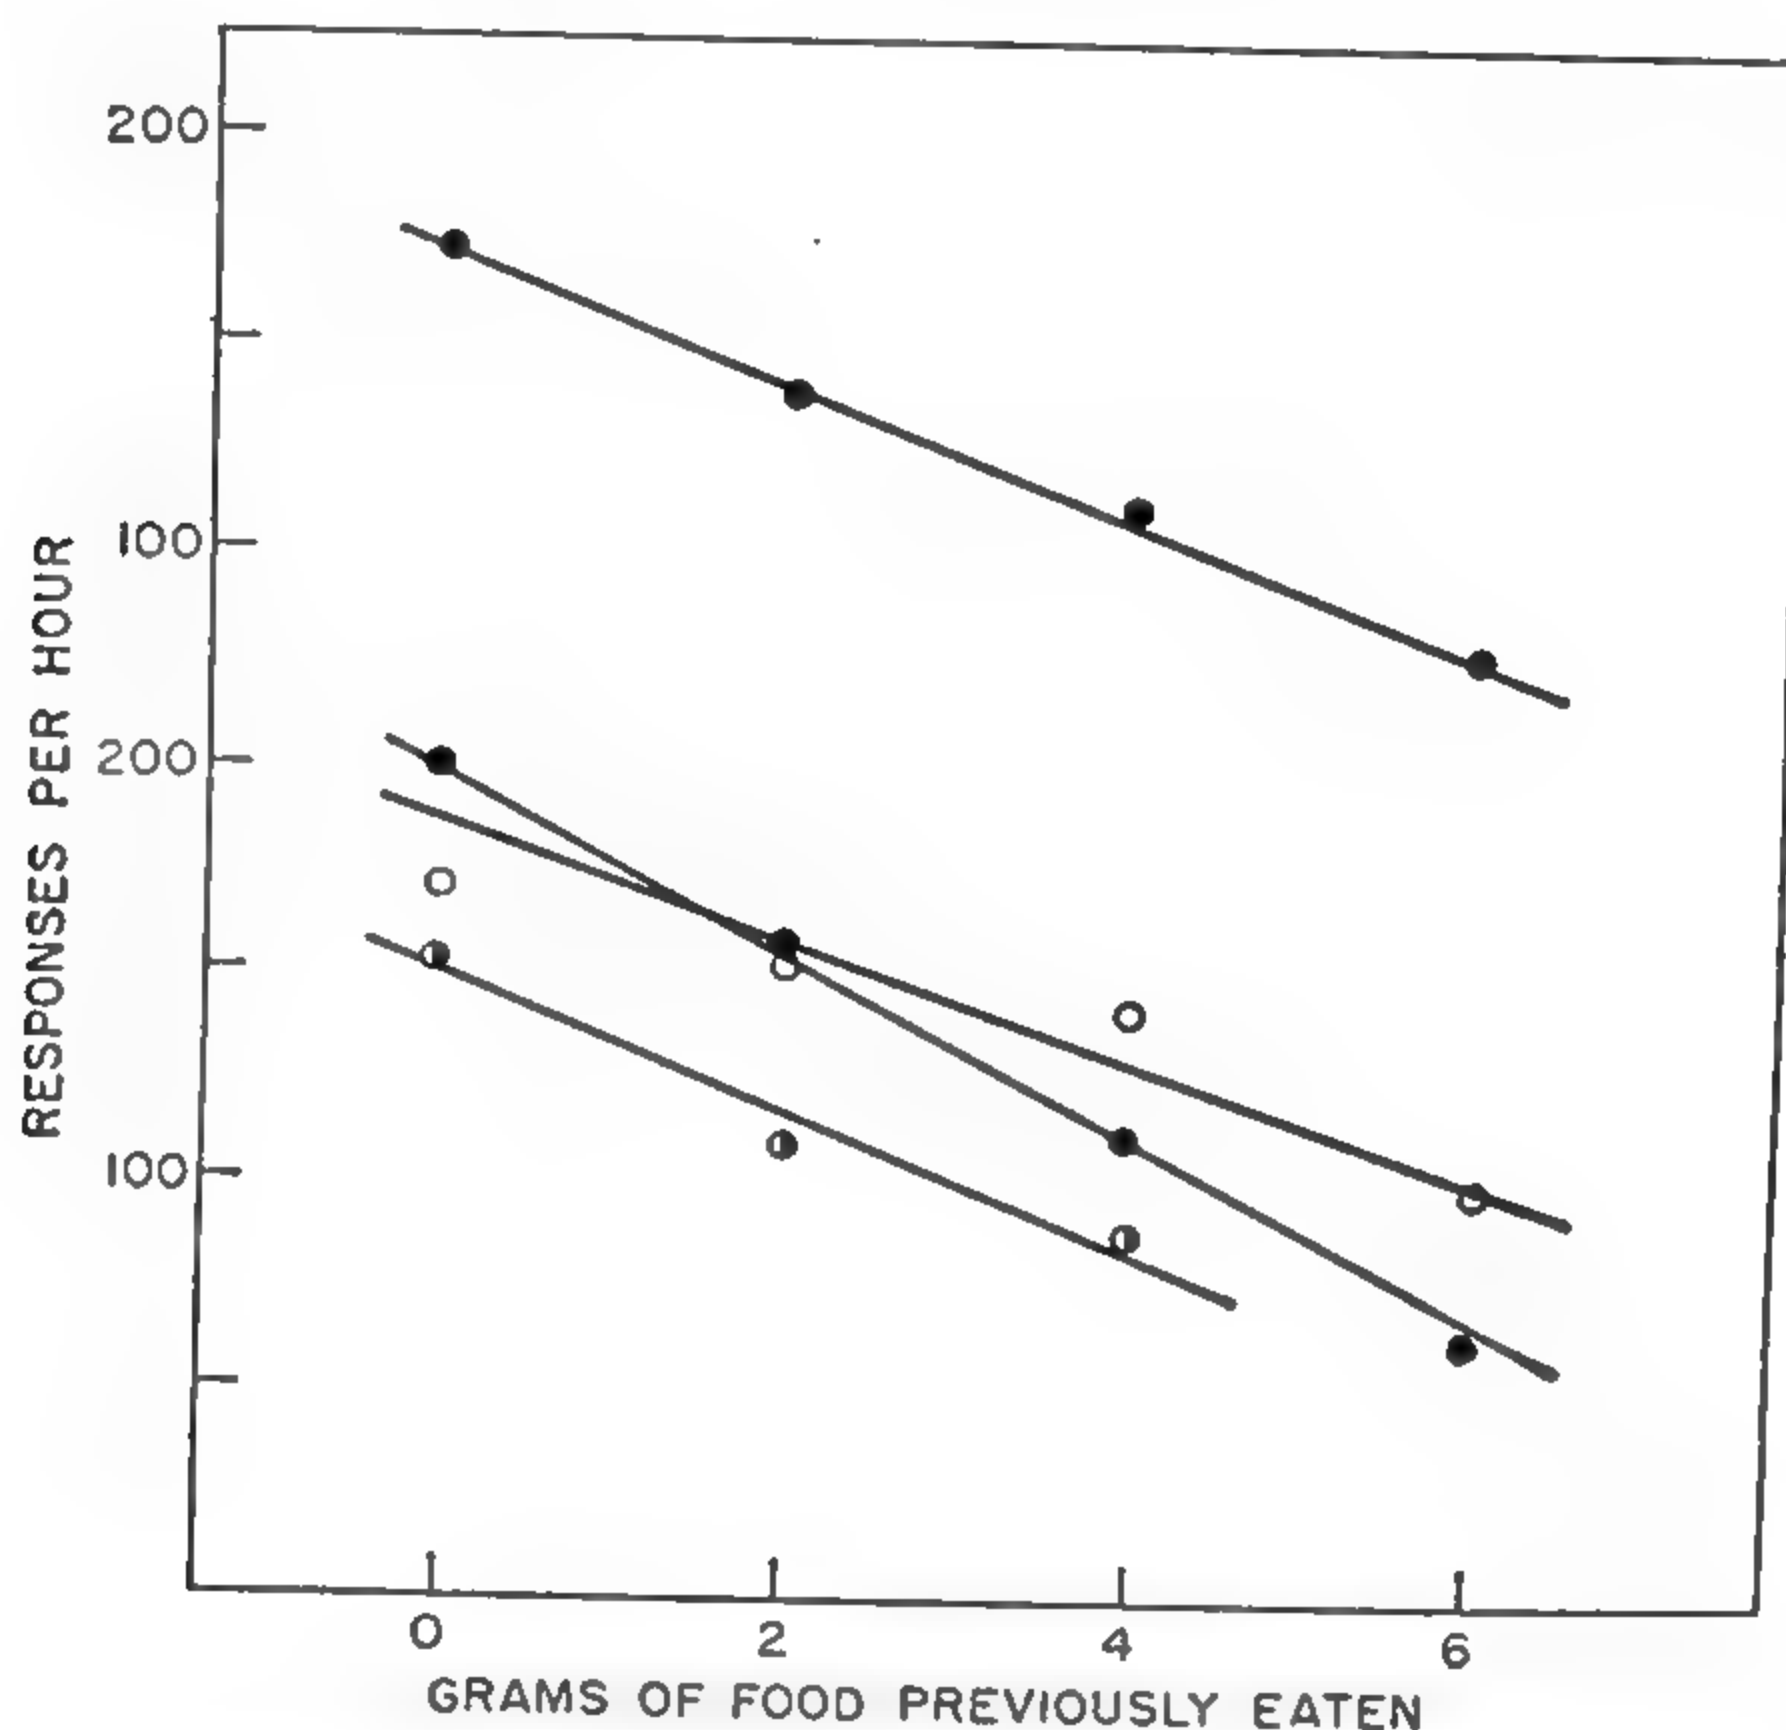


Figure 20. Rate of responding under partial reinforcement as a function of amount of food just eaten. The lower curves are based upon averages of data from four rats each. The upper curve is an average of the lower curves. (Skinner, 1938.)

Performance of learned acts and needs. As we might expect, unmotivated animals in almost any kind of learned task will make many more mistakes and in general perform less adequately than motivated ones. Even after rats have been thoroughly trained in a T maze, they will choose the alley that leads to food only about 50 per cent of the time if they are not hungry (Strange, 1950). Similarly, the rate at which rats in a Skinner box press the lever depends upon their state of hunger at the moment. Figure 20 (Skinner, 1938) illustrates this point. It shows the rate of responding after the rat has eaten various amounts of food. The more food the rat eats before the experiment, the lower its rate of response. This general effect is also

true whether or not the rat is reinforced. Figure 21 shows the rate of responding during extinction. We can see that the larger the amount of food eaten by the animals before extinction, the lower the initial rate of responding during extinction.

This increase in the level of goal-oriented behavior with an increase in the basic need seems to be generally characteristic of needs. It is true of thirst as well as hunger, and probably other needs as well. Thus, it appears that a need state and a reinforcement are both necessary for the emission of a learned instrumental response. We shall consider the implications of this for the theory of behavior shortly.

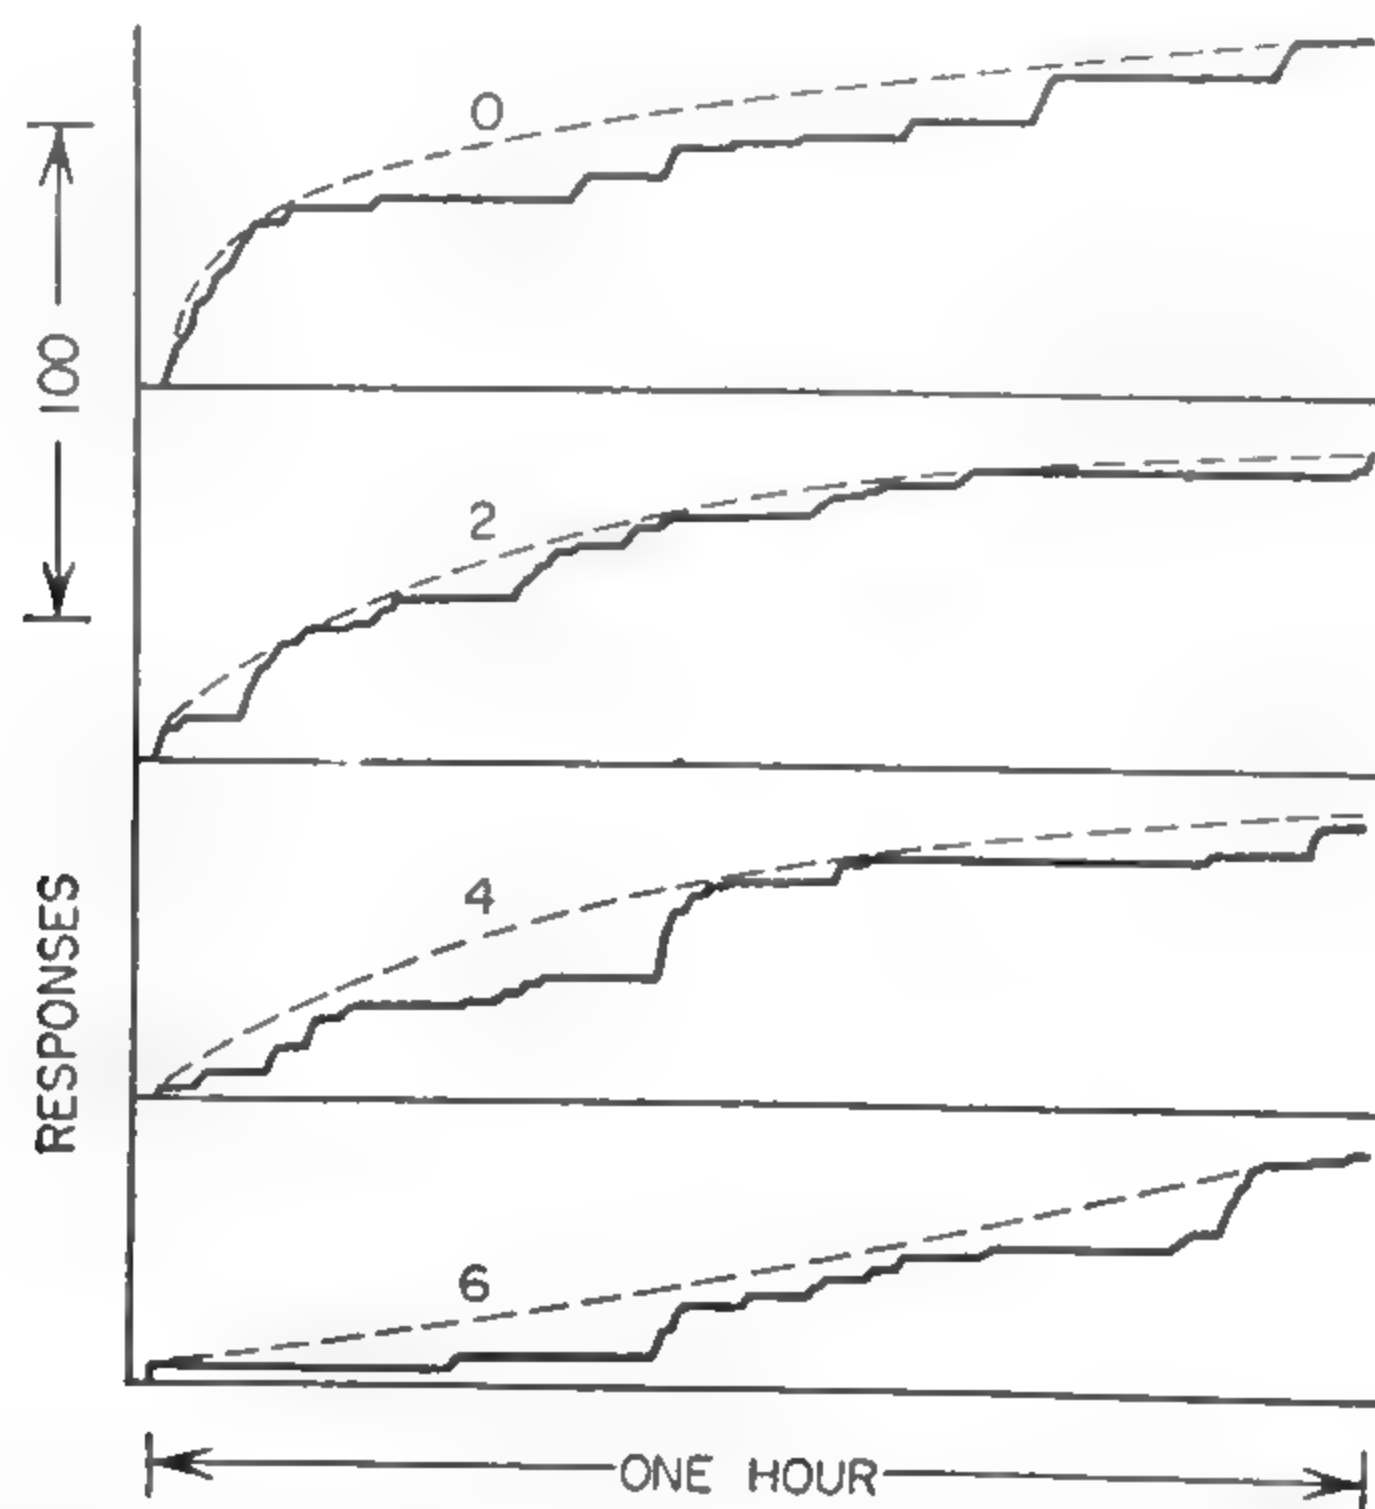


Figure 21. Examples of extinction curves for the lever-pressing habit in rats. The number shows how much food (in grams) was fed to each rat prior to extinction. Notice that the initial rate during extinction is much higher after no food than after 6 grams of food. (Skinner, 1938.)

Before we consider other topics, it is well to settle one important question about the influence of need upon learned behavior. This concerns discrimination. We know that increasing the level of need increases the tendency to perform an appropriate learned act. But does it also increase the sharpness of a learned discrimination? Suppose we have trained an animal to perform some discriminative act in which reinforcement (food) is given for pressing a lever in the presence of a light and not pressing in the absence of the light. If the level of need is then increased by starving the animal for a period of time, we find that the rate of responding in the presence

of the light is greater and that responding in the absence of light is also greater. Thus it appears that while increased need raised the absolute level of responding it did not improve the discrimination; i.e., it did not change the *ratio* of responding to the reinforced stimulus, over responding to the nonreinforced stimulus (Dinsmoor, 1952). Thus, heightened need does not appear to increase the sharpness of a discrimination, though it will increase the absolute tendency to perform the appropriate response.

Performance of learned acts and irrelevant needs. We have just seen that increasing the level of need results in greater general activity and in a greater tendency to perform acts which have in the past satisfied these needs. The question arises as to whether needs *irrelevant* to a particular pattern of behavior also tend to increase the tendency to perform that behavior. Thus, if we have trained a rat to press a lever in order to receive food, will making the rat thirsty increase the tendency to press the lever? Hull (1943), in his first important book on behavior theory, predicted, from his theory, that irrelevant needs would so affect behavior. The experimental literature does not clearly support Hull in this, however, and it will be worthwhile to see exactly how this is so.

Some experiments show that after an animal has been trained under one kind of need, increasing the level of a different kind of need increases resistance to extinction (Webb, 1949; Brandauer, 1953). This has generally been interpreted to mean that the level of an irrelevant drive does affect learned acts specific to another drive. There is, however, a serious flaw in these experiments. The animals in these experiments were trained when they were hungry. Before extinction they were made thirsty by deprivation of water, but were given complete access to food. We know, however, that rats will not eat much food when they are thirsty, and so, during extinction the animals in these experiments were probably slightly hungry as well as thirsty. This interaction between hunger and thirst has been very carefully studied, and we know that it can account for the effects in studies of irrelevant needs such as those above (Verplanck and Hayes, 1953).

In order to investigate this problem adequately, it is necessary to use drives that do not interact upon one another. Strange (1954), taking advantage of the fact that albino rats show a strong aversion to moderate light, trained a number of them to press a lever in order to turn off a light for a period of time. He then studied the effects of different levels of hunger and thirst upon the extinction of this light-aversion habit. The results clearly showed that changing the level of hunger or thirst had no effect at all upon the resistance to extinction of the light-aversion habit. Thus, it looks as

though when possible interactions between needs are controlled, there is little or no effect of irrelevant need upon habit.

Other kinds of motives. We must not assume that all of the motives of animals are based upon conditions of organic need. As we shall see later, there is some indication of an important class of learned motives. Furthermore, it now appears that there are many basic, unlearned motives that are not directly tied to some condition of organic need. After a long period in which theorists tried to reduce the motives of animals to a small list of organic need, several investigators simultaneously discovered, in different animals, that such things as visual exploration and manipulation of the environment are strong motives.

For example, Montgomery (1951, 1953), has shown that visual exploration is a strong motive in the white rat, which accounts for many of the peculiar features of maze behavior in rats. A series of studies on the rhesus monkey has clearly demonstrated that these animals are strongly motivated by curiosity and the desire to manipulate interesting objects (Harlow, 1950a; Butler, 1953; Harlow and McClearn, 1954). Not only are such motives more powerful and pervasive than anyone had thought, but it appears that they can be extremely resistant to the effects of satiation (Montgomery, 1954; Butler and Harlow, 1954). Thus they are probably very effective in producing a variety of reinforcements in learning. While it is convenient to use food or water as reinforcement in learning experiments, it is not necessary. As a matter of fact, simple curiosity about stimuli is sufficient motive to provide intrinsic reinforcement in the learning and performance of a visual discrimination in rats (Thompson and Solomon, 1954).

Motivation and learning jointly determine performance. The factors that influence behavior are classified by contemporary psychologists into two classes, associative and nonassociative. Both factors are responsible for the performance of any particular instrumental act. If an animal performs a particular act, it is highly probable that this is because (1) it has learned to do so and (2) it is motivated to do so. If motivation directed toward the particular goal or incentive associated with the act is low, the probability of occurrence of the act will be low. Likewise if habit strength for this particular act is low, the probability of occurrence of the act will be low, even though the appropriate motivation may be present.

Hull (1943) made this joint determination of response strength by motivation and habit a keystone of his systematic account of behavior, and many other theorists and experimental investigators have explored the im-

plications of the relationship. Incidentally, Hull formally specified the relationship as a multiplicative one—the effects of motivation are multiplied by the effects of learning in order to achieve the appropriate level of strength of response. In a multiplicative relationship, if either factor is absent, the product is zero, so there will be no response strength if either motivation or habit strength is zero. We cannot take this too literally, for we do not know what is meant by “zero” motivation, or even by “zero” habit strength.

Hull has shown how the joint determination of response strength by motivation and habit works by summarizing some data from other experi-

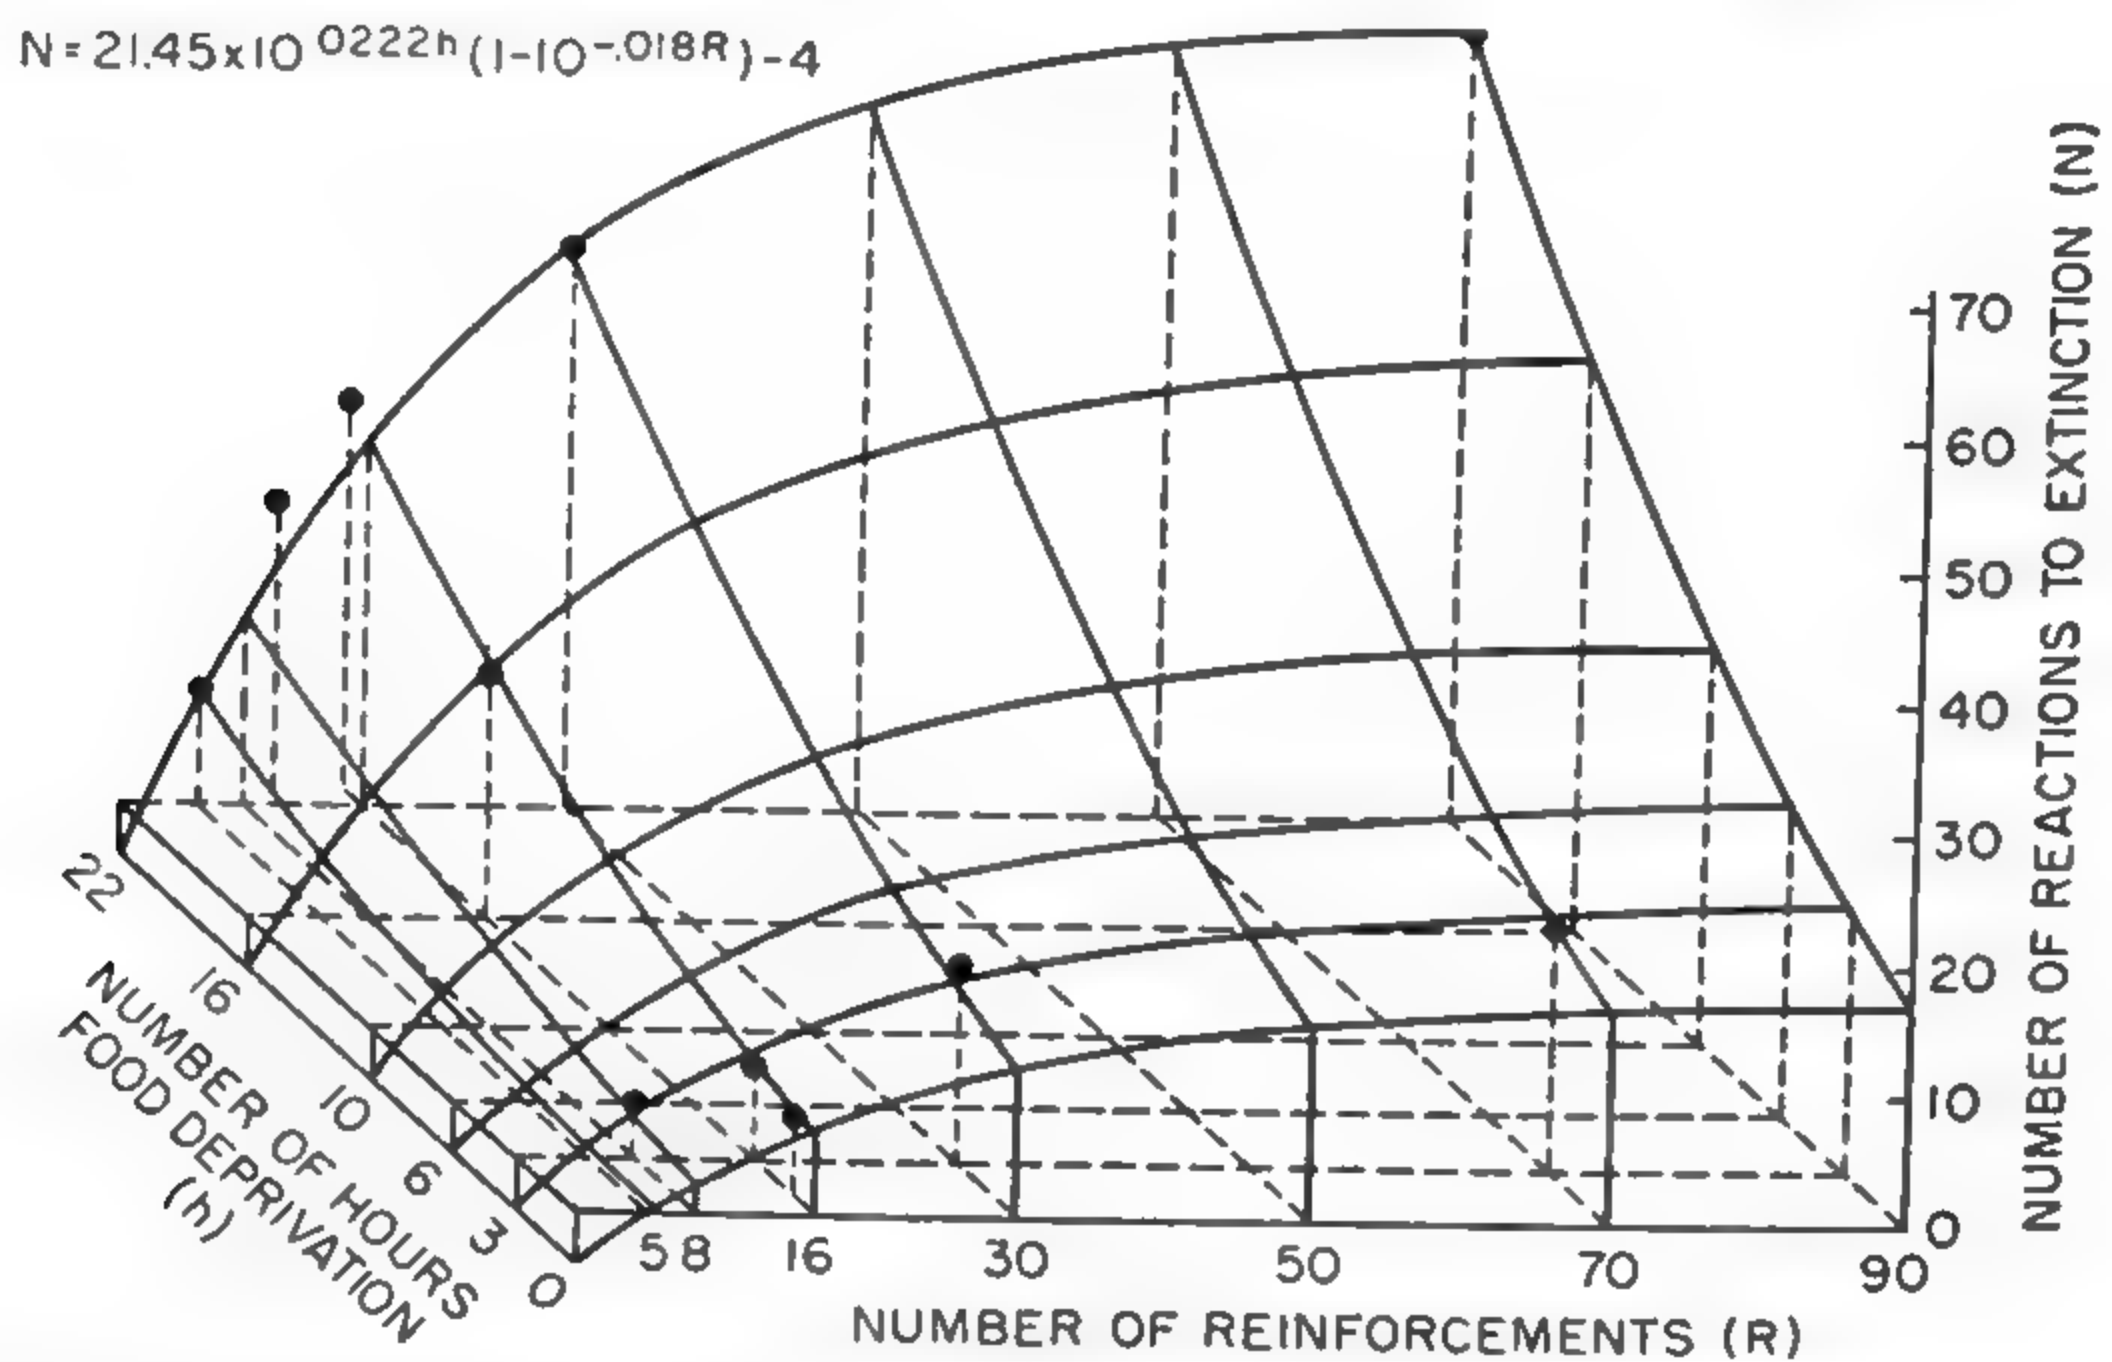


Figure 22. Resistance to extinction as a joint function of number of reinforcements and hours of deprivation. The actual obtained data are represented by the black points. The curves are derived from an equation based upon the theory of C. L. Hull. (Perin, 1942.)

menters (Perin, 1942; Williams, 1938). These investigators studied the influence of number of reinforcements (amount of training) and number of hours of food deprivation (hunger) upon resistance to extinction of a lever-pressing habit. The results can be seen in Figure 22. This figure shows something of the relationship between habit and motivation, although it seems likely that the exact nature of the relationship would be modified for other habits and other motives.

It is important to realize that this curve shows nothing of the way in which motivation might *modify* learning. It only shows how motivation and habit jointly determine the strength (measured by resistance to extinction) of a particular habit. The effect of motivation upon the *rate* of learning itself is a fundamental problem, however, and we shall deal with it next.

THE EFFECT OF MOTIVATION UPON LEARNING

In Chapter 2, we saw that one of the major conclusions to come out of the study of reinforcing stimuli is that factors relating to reinforcement—such as amount, quality, presence, or absence of reinforcement—do not intrinsically affect learning but, rather, performance. That is to say, reinforcement seems to be critical only in getting an animal to perform some act. The latent-learning experiments illustrated that with a little ingenuity it is possible to get animals to learn without specific reinforcements being associated with the act to be learned. Since motivation is simply a way of characterizing the internal state which leads an organism to respond positively to a particular reinforcement, it would seem likely that motivation would have an effect only on performance. Within limits this would seem to be true. Let us examine the experimental evidence on the question.

Latent-learning experiments. In Chapter 2 we discussed one kind of latent-learning experiment in which hungry rats were allowed to explore a maze before the introduction of a reward. In general, the results of several different experiments of this type showed, as was just mentioned, that learning was independent of the kind of specific reinforcement employed. Rats allowed to explore a maze without food reward will, under the proper conditions, show that they have learned the maze as well as rats rewarded for every trial.

The kind of latent-learning experiment important to the question of motivation and learning is a little different. The animals, instead of being motivated but not rewarded, are rewarded but unmotivated. For example, satiated rats may be allowed to explore a maze which contains a reward. If the reward is food, the animals are well fed before they are allowed in the maze. The result is that they do not seem to prefer to choose the path ending in reward; if they do find the food, they are disinterested. Since we wish to know if the rats have learned the location of food despite their indifference toward it, we make the rats hungry and then put them back in the maze. If they run the maze as well as animals receiving an equal number of reinforced trials when they are hungry, we should conclude that motivation is not necessary to learning, but only to performance.

The experiments that have been reported are somewhat more complicated than the one outlined above. Usually they have been elaborated by the addition of a second need state, so that animals are permitted to explore a maze, not just satiated with food, but thirsty. If water is present at some

point in the maze, animals readily learn to find their way to the water, but, of course, the critical problem is whether or not they will also learn the location of food. This was essentially the situation in a well-known experiment by Spence and Lippitt (1946). It turns out, incidentally, that Spence and Lippitt's rats did not apparently learn the location of food when they were thirsty and reinforced by water. As can be seen, however, the same criticism—based on the lack of independence of thirst and hunger—applies to this experiment as applied to the studies of performance and irrelevant need.

Perhaps because these latent-learning experiments have been unduly complicated, the results of different experiments have been contradictory and confusing. Many of these experiments are critically reviewed in a paper by Thistlethwaite (1951), and it is unnecessary for us to discuss them in detail. Unfortunately, however, the basic issue has been lost in many of these studies. It has been pointed out (for example, Deese, 1950) that the fundamental issue is whether some aspect of the learning function (and this can be very exactly stated as a mathematical parameter) is or is not a function of the level of motivation. Adding several different drives and rewards to the latent-learning experiments merely complicates this issue. Because so many of the latent-learning studies are thus complicated, it will be well for us to turn to another kind of experiment in order to settle this issue.

Experiments comparing effects of level of motivation. In the past few years, a large number of experiments have been designed to assess the influence of motivation upon learning. Most of these experiments have used resistance to extinction as the measure of strength of response. The experiments are usually designed so that various groups of animals are trained to press a lever in the Skinner box under various levels of motivation and then all are extinguished under a common level of motivation. If different levels of motivation produced different effects on amount of learning, resistance to extinction in the various groups should be different, even though all animals are extinguished under the same motivation, because resistance to extinction should reflect the effects of conditions of learning. Nearly all of the experiments of this sort show negative results; different levels of drive during learning have little or no effect upon resistance to extinction (Kendler, 1945; Reynolds, 1949; Strassburger, 1950). One study of this type showed slightly different results (Finan, 1940). In this study some animals were trained after only one hour of food deprivation, and others were trained after 12, 24, and 48 hours. All animals were extinguished

after 24 hours. There was slightly less resistance to extinction in the groups trained after 1 and 48 hours of food deprivation than in the 12- and 24-hour groups. While the differences found in this experiment are small, they are interesting because they suggest that the stimuli arising from various levels of hunger show stimulus generalization in much the same way as external stimuli; since the animals were all trained under moderate hunger, the center of the generalization gradient was at moderate levels.

One experiment (Carper, 1953) that has attempted to assess both the direct influence of motivation upon performance and learning in the Skinner box shows very clearly that resistance to extinction is very sensitive to motivating conditions prevailing *during* extinction but highly impervious to any differences in motivating conditions prevailing during training.

Experiments on other kinds of instrumental responses have shown more complicated results. For one thing, it seems clear that there is a *residual* effect from high motivation if animals are switched to low motivation (O'Kelly and Heyer, 1948; Deese and Carpenter, 1951). That is to say, animals trained under high motivation and then switched to low motivation will perform better than those always tested under low motivation. As in the studies of resistance to extinction in the Skinner box, however, animals trained under low motivation and then switched to high motivation do as well as those tested under high motivation all along, provided the habit is a simple one (Deese and Carpenter, 1951). Surprisingly enough, even in so complicated a habit as that required in learning a 10-unit maze, animals trained under low motivation do as well as animals trained under high motivation when tested under high motivation (Hillman, Hunter, and Kimble, 1953).

This is surprising because, as we have mentioned, the independence of learning from kind and quality of reinforcement held only for simple habits. As Seward (1951) has pointed out, a reinforcer does not strengthen a habit so much as it provides an incentive for its additional use. In complex habits the problem is complicated by the fact that a poorer quality of reinforcement provides incentive for only indifferent performance. Since animals learn only what they do or experience, if there is little or no incentive to be active, animals will not learn much. The same argument, of course, applies to need or motivation; if an animal is to learn, it must be motivated at least enough to be goaded into activity.

This does not mean that motivation and reinforcement intrinsically affect the rate of learning, but merely that if these things are missing animals will not be induced to perform complicated acts. In the case of the experiment

on maze learning outlined above, curiosity probably helped the unmotivated rats to explore the maze even in the absence of food or hunger. This motivation was weaker than hunger, consequently their performance in the maze was poorer than that of the motivated rats. It was, however, good enough to permit the unmotivated rats to learn the correct choices as well as did the motivated rats.

Thus we may conclude that for relatively simple habits, learning under low motivation is about the same as learning under high motivation. Learning of complex habits is modified by the influence of level of motivation on performance. Therefore it looks as though the associative mechanisms of learning are largely independent of the motivation-reinforcement factors that control behavior. These associative mechanisms are only incidentally brought to play in most learning situations because the combination of motivation and the reinforcement specific to that motivation induces the organism to act. Campbell and Kraeling (1954) point out that in many situations animals may learn responses that are qualitatively *different* when they learn under different levels of drive.

SECONDARY MOTIVATION AND LEARNED INCENTIVES

For a long time many psychologists took the view that there were but a handful of unlearned motives, most of which, if not all, were biological needs. If one takes this view it is difficult to describe all the elaborate motives of adult human behavior unless great emphasis is placed upon the development of learned motives. All of the complicated motives not directly and obviously connected with biological motives have been considered to be derived from biological needs by a process of learning.

Although today most psychologists grant an important role to learned motives (or at least to learned incentives), the prevailing opinion seems to be that the view stated above is too simple. We are less sure today that all *unlearned* motives are exhausted by a simple catalogue of biological needs. As we examine the behavior of animals and men in nature, it is difficult to untangle motives that seem to be learned from those that are not. It is possible that the actions of all motives, even those directly rooted in biological needs, are modified by learning and that there are some important motives that are simply the result of associative learning. Even though we are no longer convinced that all secondary motives are entirely learned, learning does play an important role in the elaboration of motives, and it will be worthwhile to examine some of the mechanisms of learning in motivation.

Secondary Drives

We saw in the chapter on reinforcement that animals sometimes come to behave toward a hitherto neutral stimulus situation as if it had incentive value. Thus, if a chimpanzee, as the result of taking part in an experiment, comes to behave as if it were fascinated with poker chips, or if a rat comes to behave as if it were mortally afraid of a flashing light, very likely we shall say that the poker chips and the flashing light have acquired secondary reinforcement value. Sometimes it has been assumed that the development of the incentive, or reinforcement, value for the neutral stimulus has been accompanied by the development of an acquired drive. There are some difficulties with this notion, as we shall see later. For the moment, however, let us examine some of these experiments on acquired or token rewards.

Token-reward studies. An interesting variation in the secondary-reinforcement studies is found in the studies of token rewards. The general idea in such studies is to choose some primary drive, such as hunger, and some token, such as a poker chip, that can be used by an animal to "purchase" food. It is not too difficult to train chimpanzees to learn to use such tokens in food-vending machines as well as to discriminate among the various colored chips according to value (Wolfe, 1936). Furthermore, the chimpanzees can be taught new habits in order to obtain the chips.

It is possible to get chimpanzees to accumulate such tokens for a while without being paid off in food (Cowles, 1937). Therefore it would appear that one could characterize the tendency of the apes to work for and accumulate poker chips as a learned drive or motive. Several theorists, however, have shown that this is a gratuitous notion (Brown, 1953; Seward, 1953). Brown points out that the pressure to assume a learned drive for poker chips would have been much greater if the apes had been either willing to work for the chips when they were not hungry or more willing to work for the chips when they could not be exchanged for food. Thus he implies that the chimpanzee's token-reward habit is based upon secondary reinforcement of a hunger motive, not upon a *new* drive for poker chips.

There are few studies in the experimental literature that suggest that animals do develop new drives by association with the reinforcement of an old one. The results of some studies are hard to account for without the assumption of learned drives; perhaps the best known of these experiments are those of Anderson (1941a; 1941b) on what he called "externalization of drive." In these, rats were required to learn a completely new maze solely on the basis of secondary reinforcement. These rats would learn the new

maze whether or not they were reinforced with food, or (and this is the important point) whether or not they were either deprived of food or satiated. The only requirement for learning to take place was that the new maze contain cues which had been associated with food in the past. Anderson characterized this kind of learning as "externalization of drive" because it showed that a motive had become established to an external object rather than an internal state of need.

Thus it seems possible that rats (and perhaps men) can acquire motives for previously indifferent goal objects. There are, however, very few experiments on this sort of thing, and therefore there are a lot of problems unsolved. What would happen, for example, if some new cues were associated with the reinforcement not of one drive, but of several drives? Would the resulting "learned" drive be more stable, less resistant to extinction, etc.? Would a partial reinforcement schedule during training make any difference in the strength of the learned drive? These are important questions, and the evidence we now have on secondary drive is not very adequate.¹ If we are able to make use of this concept in explaining more complicated problems in human behavior, our extrapolations would make more sense and perhaps be easier to assess if we had better experimental evidence in lower animals. It is still possible, for example, that Anderson's "externalization of drive" effect represents merely the arousal of a native curiosity drive. These are things that should be examined more thoroughly.

Fear as a drive. We shall have to consider in detail in the next chapter the question of fear as a learned response. Ordinarily fear is considered to be an emotion, but many psychologists have pointed out that it can have motivating properties. It is responsible for variations in the strength of certain instrumental acts, and reduction or elimination of it is apparently reinforcing. Since fear is frequently learned (though not always so, contrary to an impression one would get from reading certain parts of the current psychological literature), it is an important example of a learned drive. Unlike secondary or learned motives built upon hunger, thirst, or some other similar organic condition, learned fears are relatively easy to establish and examine in the laboratory.

Miller (1948, 1951) presents us with the classical example of fear as a learned drive. A rat is introduced into an apparatus like that illustrated in Figure 23. There are two compartments in the apparatus, a white one with an electrically wired grid as the floor and a black one with a smooth solid

¹ Wike and McNamera (1955) have made a start at answering these questions. The results of their experiments show that the answers obtained will not be simple.

floor. There is a door between the two compartments; the rat can open this door either by pressing a lever to the side of the white compartment or by turning a wheel at the end of the white compartment. First of all the rat is shocked in the white compartment and allowed to escape into the black compartment. This teaches the rat that the white compartment is to be feared and that the black compartment is safe. In order to test for a learned drive, however, it is necessary to find out whether the rat will learn a new

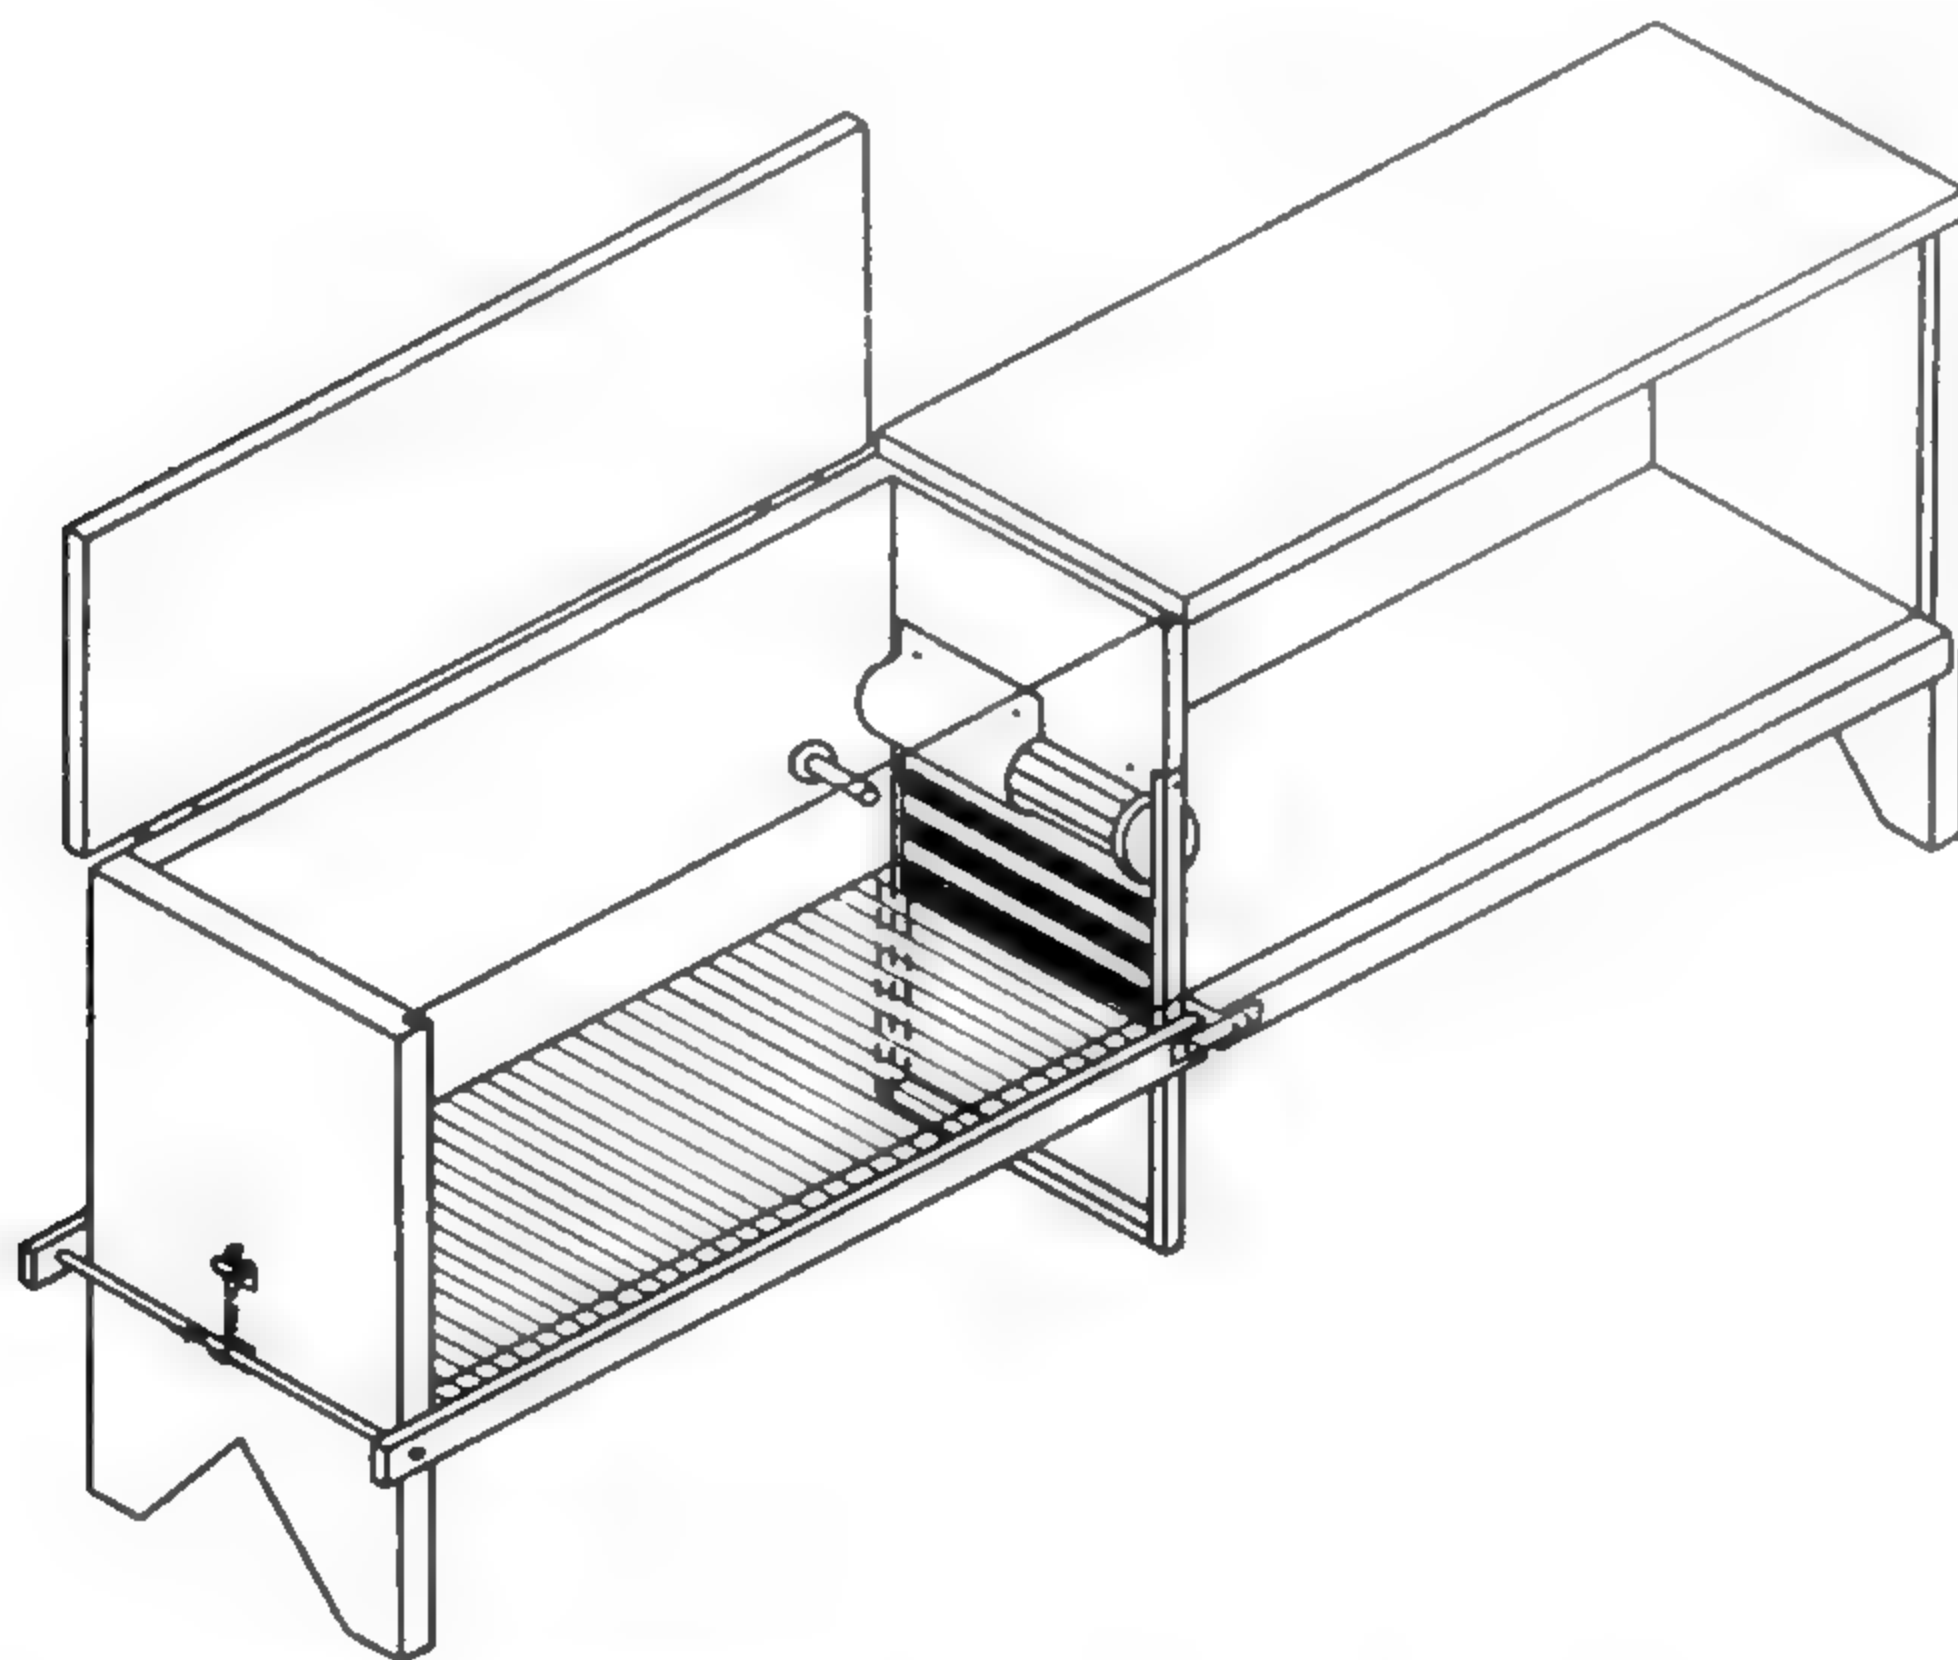


Figure 23. Apparatus for demonstration of fear as a learned drive. One compartment is white and one is black. The door between the compartments can be opened either by pressing a lever or by turning a wheel. The rat learns to escape one or the other compartment by the association of electric shock-produced pain and the stimuli from that compartment. (Miller, 1948.)

response to escape from the white compartment *without further electric shocks*. Miller showed that rats could learn to turn the wheel, with escape into the black compartment as reinforcement, much as any other instrumental response is learned. Figure 24 shows that each time the rat is placed in the white compartment it takes less time for the rat to get around to turning the wheel.

Thus it seems clear that acquired fear can serve as a motivation for instrumental activity and that reduction or elimination of this fear can serve as reinforcement. The importance of acquired fear in emotion and conflict is dealt with in the next chapter.

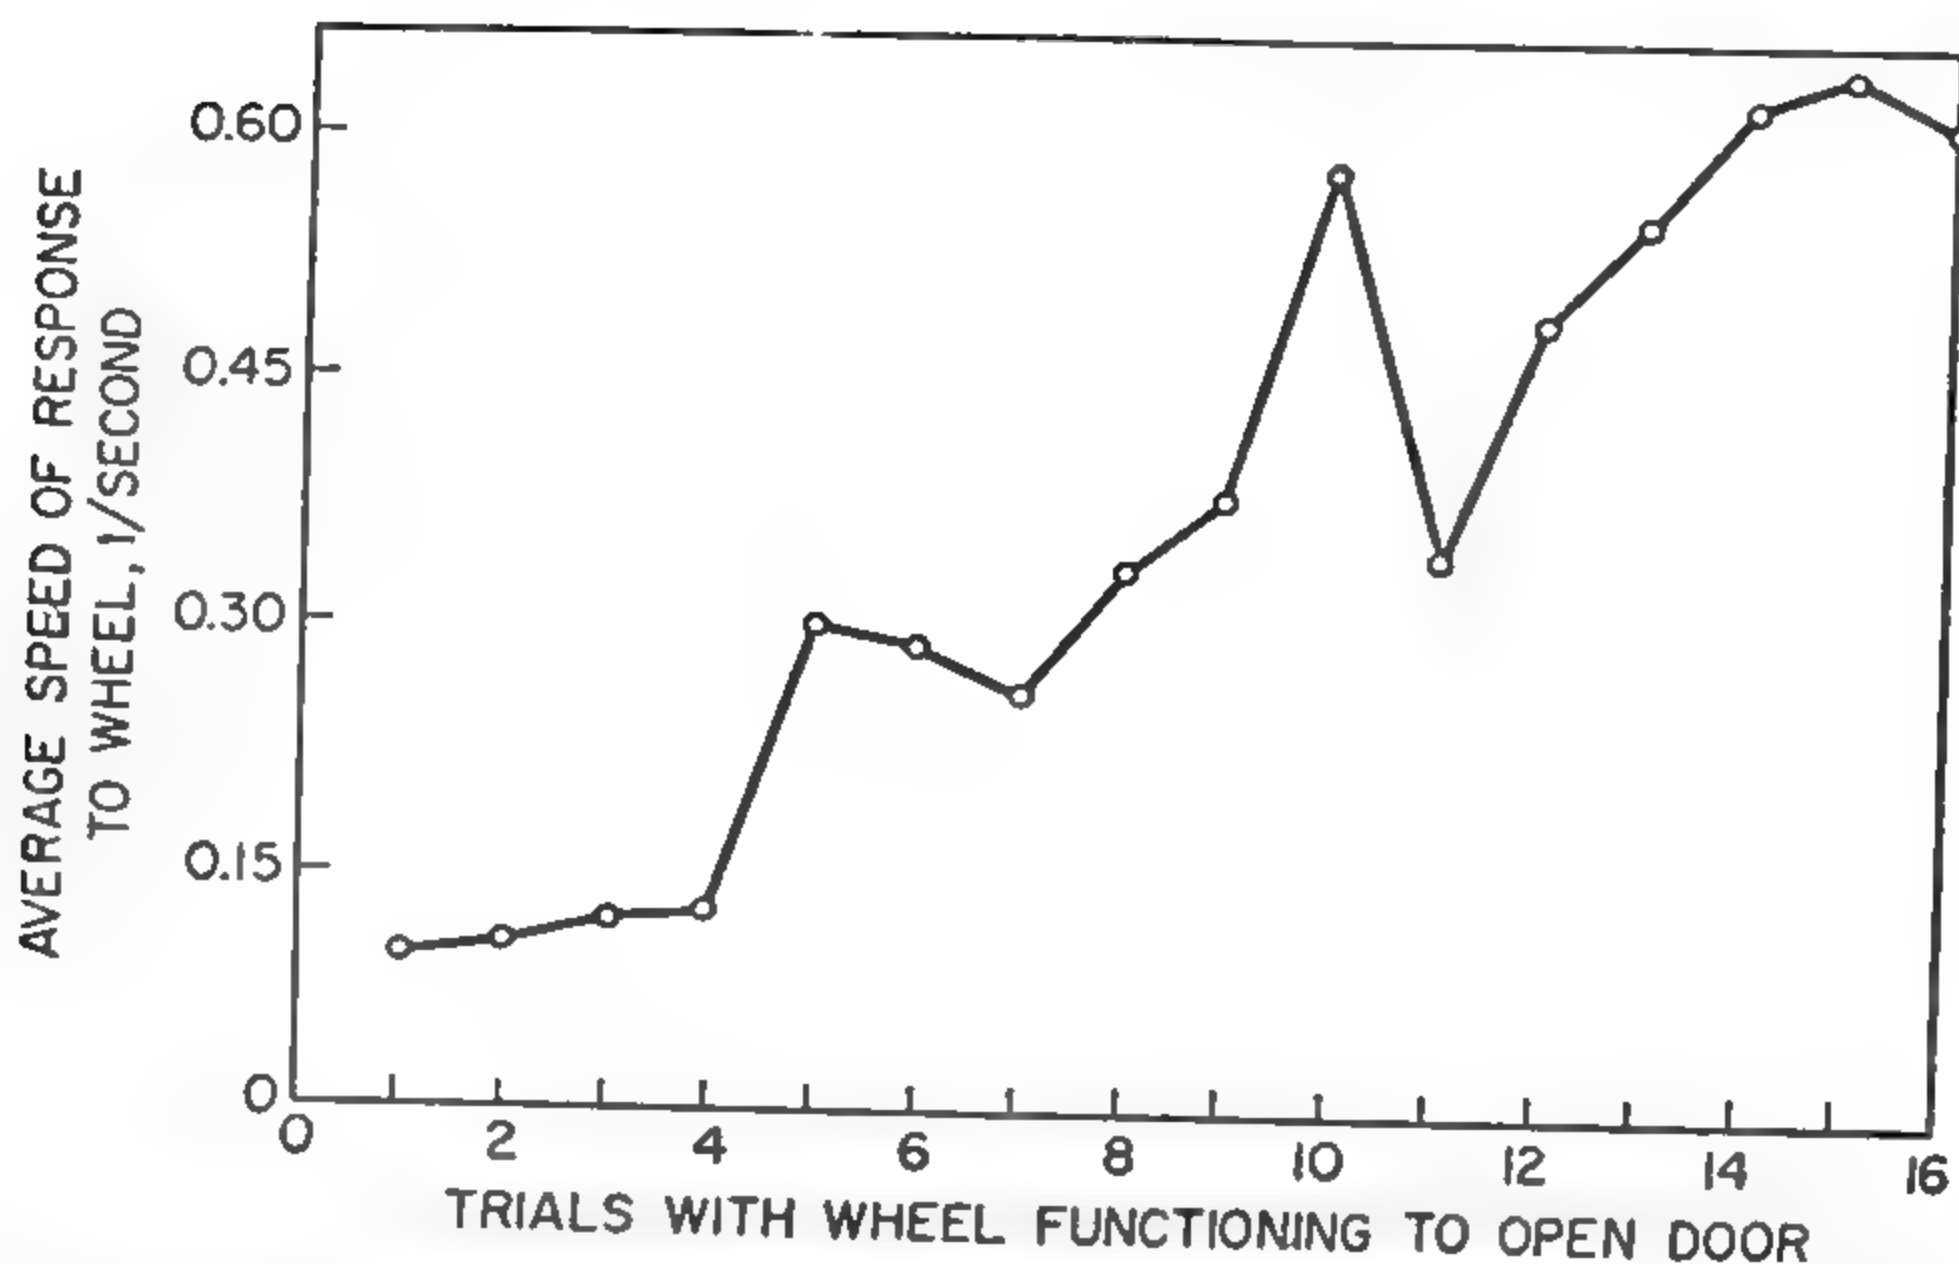


Figure 24. Learning to turn the wheel with escape from the shock compartment as a motive. The rat is shocked in the white compartment, from which it can escape by an open door. On subsequent nonshock trials, the rat must learn to turn the wheel, which opens the door, in order to escape from the compartment in which it had been shocked. (Miller, 1948.)

Theories of Secondary Drives

Acquired fear, as outlined in the experiment presented above, is fairly easy to produce. Indeed, it is sufficiently simple and reliable that it may be used in a laboratory or demonstration in elementary psychology. Learned drives based upon reward are very difficult to demonstrate, however, and, indeed, Myers and Miller (1954) present strong evidence for the existence of artifacts in the experiments purporting to demonstrate them. They suggest that these experiments (like Anderson's, mentioned above) are probably influenced by exploratory drives or curiosity.

Thus there seems to be a real difference in the ease with which secondary appetitive drives and secondary aversive drives are produced. In fact, there is not yet an unequivocal demonstration of the development of a secondary appetitive drive in the literature. Several reasons have been advanced for this difference, and it has become a matter of importance in theories of behavior.

Miller (1951) argues that it may have to do with the relative abruptness of onset of the primary drive. Aversive drives can be readily aroused by noxious stimuli (such as electric shock), so it is relatively easy to pair a particular stimulus condition with the arousal (and reduction) of the drive. Appetitive drives develop slowly; it takes time to become hungry. Therefore hunger, as a drive, cannot be neatly paired with a particular stimulus condition. Furthermore its reduction (by eating) is probably a more gradual and complicated affair.

Seward (1953) advances an argument also based on a distinction between primary drives that have sudden and those that have gradual onsets. He suggests that drives with sudden onsets may be conditioned directly to neutral stimuli (and thus give rise to fear in the case of pain as a primary drive). Drives with gradual onsets, however, are learned because of the obstruction or blocking of a learned response.

This brings us to the views of Mowrer (1953). He states that any stimulus that happens to occur with some *increase in drive strength* produces a capacity in the stimulus for eliciting a token (learned) drive. This token drive is emotional in character and akin to fear (in the case of hunger, presumably it would be fear of hunger, if it were possible to pair an increase in hunger with some simple stimulus). A token-drive decrement is reinforcing, and thus whenever a response brings about reduction of fear, this response is increased in strength.

This particular notion leads to the view that fear, or at least some kind of general avoidance, is characteristic of all secondary drives. Thus, it would seem that if the less obviously biological motivation of human life is to be explained as secondary drive, we are driven principally by fear or at least by something to be avoided. This, of course, is an unwarranted extrapolation, and to say the least, does not do justice to Mowrer's views. It does, however, suggest the more temperate and perhaps important view that much of human motivation could be conceived of as learned fear or anxiety. Perhaps one trouble with this view is the implication, which is hard to escape, that fear or anxiety is necessarily unpleasant or to be avoided. Yet it is a commonplace that human beings will deliberately place themselves in fear- or anxiety-evoking situations (such as dangerous speeding or riding a roller coaster), and it seems gratuitous to assume that people do this simply because it feels so good when they stop. It is very possible that fear or anxiety has attractive as well as negative qualities.

These matters we have been discussing are important to the theory of human motivation, and no doubt the next few years will see intensive activity in trying to fit them to problems in human behavior. At present, however, they stand as hypotheses derived from animal experiments, which have the potential of providing us with some reliable explanations of human behavior. We cannot force them to do more at present.

MOTIVE AND INCENTIVE CONDITIONS IN HUMAN LEARNING

Discussion of the nature of human motivation has probably covered more paper than any other psychological topic. We must treat this vast liter-

ature rather cursorily, for the simple reason that not very much of it is relevant to the question of the relationship between motivation and learning. Some of it is, however, and a few introductory words about theories of the origin of human motives are necessary in order to set the stage for the empirical and experimental studies we shall discuss in detail.

The Origin of Human Motives

A good deal of the theorizing about human motives has concerned the problem of their origins. An older view was that human motivation sprang from a number of instincts (the number varying with each theorist) common to all human beings and exhibited by them in varying degrees. While the multiple instinct view gives some signs of recovering from the extreme neglect it has suffered for the past thirty years, it is still not an influential notion. Today, the predominant view is that human motivation stems from either a small number of basic urges or even one basic urge, and that all of our concern for family, money, prestige, social status, security, etc., comes from an elaboration of this basic urge or urges through experience and maturation. This seems to be the fundamental viewpoint of psychoanalytic theorists, and of those theorists who would see human motivation built up through a process of secondary reinforcement and of secondary-drive learning.

Both the psychoanalytic viewpoint and the secondary-drive viewpoint seem to agree that human motives go through a developmental process. There is a great difference between these theories, of course, in emphasis, method, and implications, but their major difference as far as we are concerned here is that the psychoanalytic view would put greater stress on maturation and the secondary-drive view, on learning.

Both of these theories, however, seem to agree that goal relationships have to be learned. In other words, most of the rewards and punishments in human life are not innately rewarding or punishing, but develop these characteristics through learning. Thus the child's motives change and elaborate as it comes into various relationships with its parents and siblings.

Most psychologists are convinced that the notions of human motivation given to us by psychoanalytic or secondary-drive theory are too simple, but unfortunately there is no other systematic theory that approaches these in completeness and inventiveness. Thus, while we agree that they are probably partially right—human motivations do change through experience and maturation—we reserve the opinion that not all of the complex motives of human life can be demonstrably derived by either of these theories.

Experimental Study of Human Motivation

Experimental approach. Motivation is an invented concept used to describe certain important things about animal and human behavior; it is not a fact of the world or an experimental variable that can be directly controlled and manipulated. We associate this concept of motivation with a number of variables that we can manipulate directly, however, and these include such things as depriving animals of food for periods of time or changing their bodily balance by the introduction of drugs, etc.

In the study of human motivation we are much more limited in the factors associated with motivation that we can manipulate directly. There are a few experimental studies in which the influence of food deprivation upon some aspect of human behavior has been studied (see Lazarus, et al., 1953), but these studies are few and frequently not basic to our problems. In the older literature the most frequent technique in the study of human motivation has been to take advantage of the covariation assumed between motive state and incentive value—the assumption that most people are “more motivated” for \$10 than \$1. Thus, almost all of the experimental work on human behavior cited in P. T. Young’s well-known book, *The Motivation of Behavior* (1935), is on the effect of incentives upon human behavior. It is, however, desirable to use other techniques, if for no other reason than to study the interaction between goal or incentive value and states of the individual induced by other conditions of his life.

It is extremely difficult to produce conditions in the experimental laboratory that allow us to vary systematically the factors associated with human motivation. For this reason, in recent years, there has been an increased emphasis in the study of human motivation upon what has been called response-response relationships. This means that we use some behavior of the individual to tell us how he is motivated, and then use this behavior to predict something further about an individual—perhaps how he will perform on a test of learning. Thus one of the main techniques in the study of human motivation has been the correlational approach. There are certain inherent limitations in this approach, the principal of which is the inability to eliminate extraneous variables. Thus, there is considerable argument over the merits of the response-response technique, or correlational approach, to the study of human motivation.

Studies of incentive value. There is a large literature in industrial psychology on the relationship between various incentives and output. Most of this is not particularly relevant to learning, but two examples will demon-

strate a point. Kitson (1922), in a classical study, showed that industrial workers of long experience may not necessarily be working at top efficiency. Despite the fact that the typesetters whom he studied had ample opportunity to learn their trade to perfection, the addition of a monetary incentive resulted in a great improvement in output. Similar studies (Mace, 1935) have shown that the addition of an incentive may delay or eliminate work decrement due to continuous work.

Various incentives used in industrial work create many problems of a psychological and sociological sort. While these are fundamental to an understanding of human behavior in industrial society, they are not directly important to the psychology of learning. Consequently, the techniques of incentive payment in industry will not be stressed here. Suffice it to say that the use of incentives such as piece-rate payments or special bonuses does not always work as we think it should. The motivations of workers in modern industrial society are very complicated. It is true that the human being in our civilization seems to be almost universally motivated for the acquisition of coin, but this is frequently less important than other motives. He may be motivated to earn a certain, minimum amount of money with little effort. He may be motivated by such things as security, social approval, feelings of self-respect, etc. Thus, the manipulation of money as an incentive may not always lead to the expected results.

In the school, grades, verbal rewards, and reproofs have been intensively studied. These factors probably operate less upon learning itself than upon performance, but there is little reason to doubt that they are powerful techniques in the control of performance level. The effects of specific verbal rewards and punishments are sometimes difficult to evaluate, however, because there are enormous individual differences in the effectiveness of these incentives (differences probably associated with individual differences in basic patterns of motivation).

Motivation assessed by a self-report inventory. Let us turn to the technique of assessing motivation by means of the behavior of the individual. The most commonly employed technique here is the use of self-report inventories. These consist of a series of items such as "I am sometimes troubled over what others think of me" which the individual answers by true or false alternatives. In the experimental study of motivation the largest number of investigators have made use of a self-report inventory known as the Manifest Anxiety Scale (Taylor, 1956, etc.). The theory behind this inventory is that the items in it tap an emotional responsiveness that is related to drive or motivational level. Those individuals scoring high on the inven-

tory are said to be anxious, and their anxiety state is assumed to act as a generalized motivator to behavior. Because there is some emotional basis to this kind of anxiety drive, it has been questioned whether the Manifest Anxiety Scale primarily reflects properties ordinarily associated with emotion or with motivation (Deese, Lazarus, and Keenan, 1953; Eriksen, 1954). Consequently we shall deal with many of these same problems in the chapter on emotion and learning. For the present, however, let us see how far a consideration of manifest anxiety as an indicator of motivation gets us.

At first blush the results of experimental studies are very favorable to the interpretation of this inventory as a measure of anxiety drive. There seems to be little doubt that classical conditioned responses (either the galvanic skin response, or a skeletal reflex such as the eyeblink) occur more readily in highly anxious individuals than in individuals reporting little anxiety (Taylor, 1951; Spence and Taylor, 1951; Bitterman and Holtzman, 1952; etc.). This is what one would expect if high drive had a generalized effect of increasing tendency to respond.

Beyond these studies of classical conditioning the results become less clear, however. For example, in a study of the frequency of eyeblinks (not conditioning, but simply a study of frequency of occurrence), there was no tendency for individuals high on the anxiety inventory to blink more frequently, even when under conditions of tension (Meyer, Bahrnick, and Fitts, 1953). Curiously enough, however, there was a tendency for frequency of blinks to correlate with another inventory of maladjustment.

When we turn to the study of more complicated, instrumental acts, the picture becomes even more confused. Here, very frequently high anxiety is associated with poor performance (Taylor and Spence, 1952; Farber and Spence, 1953). High drive leading to poor performance would, at first glance, seem to be contrary to the usual effects of drive, but this is not necessarily so. For if anxiety is a truly generalized drive it will activate incorrect as well as correct responses. If, as is frequently the case in complex human learning, incorrect responses are initially stronger than correct responses, it is argued that the effect of increasing drive would be to increase the difference between incorrect and correct responses. Therefore, one might reasonably expect, so the argument goes, that high anxiety would be associated with poor performance in some situations, at least initially. Furthermore, one experiment (Montague, 1953) shows that if the task to be learned gives ample opportunity for interference, subjects scoring high on the anxiety inventory do more poorly; but if the task gives little opportunity

for interfering effects of incorrect on correct responses, subjects high in anxiety do better.

Now this seems all very well, except that a direct test of the influence of drive state upon the interfering tendencies of incorrect responses does not seem to bear the analysis out. Buchwald and Yamaguchi (1955) compared the effects of habit reversal in the maze learning of rats under high and low thirst. The results clearly showed that reversal learning was more rapid under high drive, despite the fact that reversals in maze learning provide just about the optimal conditions for interference.

The fact is that a hypothesis about the effect of drive upon interfering tendencies has never been clearly stated. Taylor and Spence (1952; Taylor, 1956) say that because drive multiplies habit strength to determine response strength, increasing the drive ought therefore to increase relatively strong habits at the expense of weak habits. Thus, if incorrect habits are strong they ought to be increased more than weak habits. This assumption, however, is not at all clear from the mathematics. To be sure, multiplying two different numbers by a constant (equivalent, according to Taylor and Spence, to the effect of high drive on strong and weak habits) increases the *absolute* difference between them, but the *ratio* of the two habits remains the same. Some theorists have assumed that any choice an organism makes between two responses is based upon the ratio of their strengths (or their probability of occurrence, if the two responses exhaust the possibilities), and in this case increasing drive level would *not* differentially increase the strength of the stronger habit.

Thus, while the idea of using the Manifest Anxiety Scale as an index of motivation is an interesting one, the implications have been somewhat disappointing. In the next chapter we shall examine some of the implications of using this as well as other, similar scales as indices of emotional reactivity to stress. As we shall see, it is probable that our current concepts of emotion and motivation are closely allied to each other and that any really serious attempt to decide whether the Manifest Anxiety Scale reflects more of one than the other will await clarification of the concepts themselves.

—Another kind of anxiety inventory has been used by other investigators (Mandler and Sarason, 1952; Sarason, Mandler, and Craighill, 1952). This inventory was designed to assess the individual's reaction to situations in which his abilities are under test. Does, for example, an individual display signs of anxiety and worry when he is taking an important examination? The authors of this inventory believe that it taps both motivation for achievement and emotional responsiveness to conditions of stress. Thus in-

dividuals scoring high on this inventory are not only higher in motivation, but also, because of past learning, are likely to exhibit signs of emotional disturbance which might interfere with performance.

Generally speaking, individuals who score on the high end of this "test anxiety" inventory do more poorly in tests of learning and performance. This is particularly true in the case of tests known to be sensitive to emotional interference. In some cases, however, these individuals do better than those scoring low; they increase their scores more with practice. This is probably due both to the fact that these individuals adapt to their states of emotionality and to the fact that they are probably more highly motivated.

Thus, the studies of individual differences in the things tapped by these personality inventories suggest that in human life, there is a close relationship between many kinds of motivation and emotion. These are so intertwined with one another that it is probably impossible simply to say that higher motivation automatically means better performance. This also depends upon the kind and level of emotional reactivity.

✓*Projective tests of motivation.* Another important technique in the assessment of individual differences in human motivation lies in projective tests. The most ambitious of these has been an attempt to measure need for achievement by means of the Thematic Apperception Test (McClelland, et al., 1953). In this test, an individual's need is evaluated by means of the stories he makes up to fit pictures of individuals in various situations. In general, great reliance is placed upon the extent to which an individual shows achievement imagery. The authors of this version of the Thematic Apperception Test were able to bring impressive evidence to show that they were actually measuring individual differences in motivation.

What about the need achievement-learning and performance relationship? One experiment (Lowell, 1952) rather clearly showed that there was a difference in performance during the learning of an anagrams word test. Those individuals high in need achievement did much better. A test on another task, in which subjects were not expected to show learning or improvement, suggested that the difference in performance associated with need achievement is correlated with performance level and not with rate of learning. This is, of course, much what we should have predicted from studies of animal behavior. This relationship between need achievement and learning and performance may eventually be of practical importance, since there is evidence that need achievement scores are correlated with college grades, even when individual differences in intelligence are taken into consideration (McClelland, et al., 1953).

Thus, although the use of individual differences as a technique for investigating motivation is relatively new, we have encouraging evidence that we shall learn something by it. The use of personality inventory scores has proved to be a little disappointing, largely because they are heavily saturated with emotional factors. There are many possibilities, however, for the development of inventories relatively free from emotional components (perhaps something like the traditional study habits inventory will work with college students), and this matter is not yet closed. The need-achievement scoring of the Thematic Apperception Test gives great promise, and recently a similar projective test has been used, with encouraging results, in the study of achievement motivation (French, 1955). We do not have much data on the relationship between these tests and the question of whether learning or performance is affected by motivation, but the evidence we do have suggests that it is largely performance rather than learning that is affected by different levels of motivation in human beings. Thus, any differences that do develop in learning are probably secondarily due to the basic differences in performance level.

Incidental Learning

We find another important technique in the study of motivation in human learning in experiments on incidental learning. In these experiments the performance of individuals who practice at a task with deliberate intent to learn is compared with that of individuals who practice without this intent. Let us look at a simple example of an experiment in incidental learning.

An experiment in incidental learning. One example of incidental learning is provided in a study by Jenkins (1933). He had students, who thought they were serving as experimenters, read lists of nonsense words to other students, who were instructed to learn the words. After the learners had reached a certain criterion of performance, both the subjects and the experimenters were dismissed and told to return 24 hours later. At this time both subjects and experimenters were asked to recall the lists of words. It turned out that both the learners and the experimenters recalled many of the words; the learners, however, recalled more words than did the experimenters. Upon asking for reports from the subjects, however, Jenkins found that some of the experimenters did deliberately try to learn, but even those who did not, recalled many words.

This experiment shows that (1) even without specific instructions to learn, human beings will form a set to learn; (2) even without this specific set to learn, human beings will learn. In this connection, we may note, as

the set to learn decreased in Jenkins's subjects, the amount of material retained 24 hours later also decreased. Thus, while it appears that subjects can learn without a specific set to learn, they seem to learn more adequately when they have such a set.

Analyses of incidental learning. Incidental learning is a complicated effect, and there have been many experiments performed in an effort to untangle the basic causes of the differences between incidental and instructed subjects. Right now, of course, we are interested in the question principally from the point of view of motivation. Is the lesser motivational set induced in the incidental learners responsible for their poorer performance?

In order to answer this question we must look more closely at the experiments. Saltzman (1953), for example, asked whether or not the difference between incidental and intentional learning might be due to the difference in the orienting task required of subjects working under two different conditions. Translating Saltzman's question into the context of the experiment outlined above, we can see that even the task of being an "experimenter" might interfere with the incidental learning of words. Saltzman was able to show that if the orienting task were allowed to interfere with the performance of *intentional* learners, they actually did worse than incidental learners.

In subsequent experiments (Neimark and Saltzman, 1953; Saltzman, 1956) the problem was further explored; it turns out that if intentional learning is superior to incidental learning, this is so only under certain conditions (such as certain distributions of trials through time) and with certain orienting tasks. Some orienting tasks given to incidental learners interfere with learning and others facilitate learning as much as does intent to learn.

Further analyses of problems in incidental learning have been made by Postman and his associates. Like Saltzman, they (Postman, Adams, and Phillips, 1955) assume that the difference between incidental and intentional learning depends upon the kinds of responses people make to the material presented to them. Specifically, they thought, intent to learn would make it more likely that human learners would notice critical differences between various verbal items they were required to learn. The experimental results showed that when items to be learned were easy to discriminate, there was not much difference between incidental and intentional learning; but in those cases in which the items to be learned were unfamiliar and low in association value, there was a large difference.

Other experiments in this same series (Postman and Adams, 1956; Postman, Adams, and Bohm, 1956) further emphasize the role of intent to learn

in determining the kind of responses made by the learners and whether or not these will be appropriate to the task of learning. Thus it is fairly clear that the comparison between incidental and intentional learning is not simple, since the different effects these two conditions have upon performance will depend upon the nature of the task to be learned and the learner's orientation to the task.

Incidental learning and performance versus learning. What do the experiments on incidental learning tell us about the comparative effects of motivation upon learning and upon performance? It is clear from these experiments that *what* is learned depends upon the given set for learning. To the extent that a quantitative index of learning depends upon what is learned, the amount learned will depend upon the set for learning. Thus it is not surprising that an attempt to find an analogy in human learning to the latent learning found in maze learning of rats was not too successful (Postman and Tuma, 1954). The troublesome thing about this experiment seems to be that it is difficult or impossible to devise conditions in which subjects perform the correct responses for learning the task. It is probably unfair to push the analogy to maze learning in rats too far, but it is much as if an unmotivated rat refused to run through the maze and thus could hardly learn the maze, latently or otherwise.

Therefore, it is clear that unmotivated people sometimes do not learn very well, not because of some intrinsic connection between motivation and rate of learning but because they do not perform the responses actually to be learned. Some years ago there came to the attention of the author a youngster who had completely failed to learn the mechanics of subtraction in school. He was a bright boy, but evidently during the school hours devoted to subtraction he had hit upon the expedient of writing down numbers more or less at random in order to satisfy the superficial requirements of the classroom. Thus, this unmotivated youngster failed to learn subtraction for the simple reason that he would not go through the activity necessary to the learning of subtraction. In a practical way, motivation may be important to learning simply because it is important to the performance of the right responses.

CHAPTER 6

EMOTION AND LEARNING

It is apparent to everyone that learning plays an important part in emotional and affective behavior. Most of our likes and dislikes are learned. We all know that many, if not most, fears are acquired. We know about the importance of experience in the development of personality disturbances. In this chapter we shall not consider all of these things, but rather, look at some of the basic principles we must use when we apply the psychology of learning to problems in emotion.

In applying the psychology of learning to problems in emotion, some new and important issues arise, and it is these issues to which we shall largely direct our attention. The problem of emotional learning does not involve merely the same principles we have discussed already; it involves new principles, as well as new ways of considering those we have already looked at. Before looking at the application of the psychology of learning to emotion, however, one basic point about the concept of emotion needs to be mentioned.

Even the most superficial examination of the biological and psychological characteristics of emotion shows that it is related to the concept of motivation. Motivation, as it is most often used by psychologists, emphasizes the activating role of central mechanisms in behavior. It has been apparent for a long time, however, that the factors associated with the activation of behavior are also associated with feelings of pleasantness and unpleasantness, as well as with expressive behavior in rage or fear (Lindsley, 1951). Recent anatomical and behavioral work has shown the intimate interaction between the central neural mechanisms associated with motivation and those associated with emotion (see Stellar, 1954).

It is certainly apparent that some kinds of emotional activity accompany powerful motivation. We saw in the last chapter, for example, that a learned fear can motivate rats to learn a new instrumental act. The fear is clearly an emotional state, but it is also a motivational state. It turns out that a large portion of the experimental literature on the relationship between emotion and learning is devoted to an analysis of the motivating properties of such powerful, unpleasant emotions as fear.

Before we consider the motivating aspects of emotion, however, let us see how emotions themselves are learned at the most elementary level. In order to do this, we shall examine some studies of emotional learning in animals and very young children.

CONDITIONING AND EMOTIONS

An important point about the psychology of emotions is that it has always been associated with the activity of the autonomic nervous system—that part of the nervous system controlling the actions of the heart, glands, and smooth muscles. The activity of the autonomic nervous system is not voluntary, and the elements of emotional expression autonomically controlled are automatic and reflex-like in character. Thus, acceleration of the heartbeat, sweating, blushing, etc., are all things that occur more or less automatically, given the presence of some emotion-producing condition.

This involuntary characteristic of elements of emotional expression is important, because it suggests that the conditioning of these responses is classical rather than instrumental. In the first chapter, we saw that in classical conditioning responses were elicited directly by stimuli. An unconditioned stimulus elicits a response and this response can be transferred to a hitherto ineffective stimulus simply by pairing that stimulus with the unconditioned stimulus. In instrumental conditioning, the response usually occurs spontaneously—there is no stimulus that seems to elicit it—and this response is learned (occurs more frequently) when it is followed by some reinforcement or reward. Thus, since emotional activity seems to be elicited by some more or less well-defined stimulus, it can be conditioned by the classical conditioning technique.

Because of the reflex-like characteristics of emotional behavior, the first applications of Pavlov's technique in American laboratories were to the conditioning of emotional behavior. Let us look at some of these early attempts to apply the classical conditioning technique to learning emotional behavior.

Classical Conditioning of Emotions

(*The Watson and Raynor study.* One of the best known of all these studies is by Watson and Raynor (1920). These investigators were interested in the development of fears in infants, and they made a long series of observations on a number of infants. One baby boy, who was about one year old at the time, was the subject of their conditioning study. Previously, they had discovered that a number of stimuli which one might ordinarily think would arouse fear were not fear producing for this youngster. White rats, dogs, masks, burning newspapers—none of these produced fear. One stimulus that did produce a startle reaction followed by crying and other signs of emotional activity was a loud sound. Consequently, they decided to employ the loud sound as an unconditioned stimulus for emotional activity.

In the actual conditioning procedure they paired the sight of a white rat (previously not fear producing) with the occurrence of the sound. A very few pairings were sufficient to elicit a full-scale emotional response simply by presenting the white rat alone. Interestingly enough, this conditioned fear generalized to other animals and to other furry objects, such as a wad of cotton.)

This was a rather clear and dramatic demonstration of the possibilities of conditioned emotions, as well as a demonstration of the importance of such principles as stimulus generalization in the extension of the conditioned reaction. Watson and Raynor thought that this experiment could serve as a model for the way in which the complex and irrational fears of daily life might arise, and they placed much emphasis upon the importance of conditioned emotional reactions in the child's personality development.

In the intervening years, however, there have been enough studies of the developmental processes in young children so that we know Watson and Raynor's conclusions were a bit overdrawn. We now know that the specific stimuli for fear and other emotional reactions depend to a considerable extent upon maturation and the particular developmental stage of the child. In part, the high specificity Watson and Raynor attributed to the emotional reactivity of the child was due to his age and maturation level.

Other studies of emotional conditioning in children. One important point from the work of Watson and Raynor stands up well, however. We know from subsequent investigations that emotional reactions can be conditioned by the classical technique. Jones (1931), for example, was able to show much the same pattern of conditioning in a fifteen-month infant by

pairing an electric shock and a buzzer. Jones studied the extinction of this conditioned response and noticed that overt reactions to the conditioned stimulus extinguished much more rapidly than signs of autonomic activity. Jones also established the existence of such phenomena characteristic of classical conditioning as external inhibition and spontaneous recovery.

In another study (Jones, 1930), the same investigator conditioned the galvanic skin reaction in children by pairing a mild electric shock with a buzzer. The children in this study ranged in age from three to nine months, so that the investigator was able to establish the fact of conditioned emotional reactions early in infancy. Jones concluded that his studies rather clearly established the applicability of the classical Pavlovian conditioning technique to the modification of emotion in children.

The importance of emotional conditioning. It turns out that these simple examples of conditioned emotional reactions are extremely important in the theory of learning, particularly as the theory applies to situations in which emotional disturbances have occurred. The essential point of these studies is that emotional activity is conditioned by the classical technique. Since the emotional reaction is aroused originally by an unconditioned stimulus, there is no need to control it through reward. It is behavior that is, by and large, beyond voluntary control. Few of us can directly elicit emotional activity in ourselves in the same sense that we can open doors and recite the pledge to the flag. A few people can "voluntarily" cry, etc., but there is a strong belief that this is an indirect kind of control. That is to say, it may be possible to produce voluntary crying by thinking about something likely to arouse emotional activity. The crying is indirectly aroused by the presence of some internal (perhaps conditioned) stimulus. It is, incidentally, possible to make use of "lie detection" techniques simply because emotional reactions usually beyond voluntary control can be aroused by external stimuli.

Thus, the conditioning of emotional reactions seems to be largely a process of association. Organisms learn to become emotional in new situations simply because these situations occur when they are emotional for other reasons. At the beginning of this chapter, however, we pointed out that emotions seem to have some of the properties of motivating states such as hunger or thirst. Where, then, does this aspect of emotional reactivity fit in? Clearly not in the transference of emotional reactions themselves from an unconditioned to a conditioned stimulus, since this merely involves the association of the response elicited by both these stimuli. If the emotional reaction is itself motivating, then *reduction* of it ought to serve as reinforce-

ment for instrumental activity. This gives rise to some very important questions, which we can best examine by looking at the problem of avoidance conditioning.

Avoidance Conditioning and the Two-factor Theory

Since experiments in avoidance conditioning have been at the root of most notions that have tried to describe the role of the motivating properties of emotions in learning, we shall do well to begin our discussion by looking at this kind of conditioning.

An example of avoidance conditioning. Avoidance conditioning causes much difficulty because it looks very much like classical conditioning, although, as we shall see, the resemblance is misleading. A simple example will illustrate this point.

A dog is trained to stand perfectly still in an experimental apparatus which permits the careful recording of the movements of all of its legs. One of its rear legs rests upon an electrode that completes a circuit through the dog, so that it can be shocked through this leg. If the dog lifts its paw, of course, the shock is terminated. In the conditioning procedure, a buzzer sounds about two seconds before a brief shock is applied to the dog's leg. Soon after the conditioning procedure has begun, there are obvious signs of emotional disturbance elicited by the buzzer, and this makes it look very much like a simple example of classical conditioning. After a bit, however, the dog begins to lift its paw when the buzzer comes on. If the paw is kept up until the buzzer stops, the dog will not get shocked. After quite a few training trials, the dog lifts its paw to every sound of the buzzer, and thus never gets shocked. Therefore, it is quite clear that the paw-lifting response is *instrumental*; it prevents the dog from being shocked.

This experiment starts out as an example of classical conditioning, but it clearly ends up with the characteristics of instrumental conditioning, for the response permits the dog to avoid the electric shock.

This experiment could have been turned into an example of a "pure" classical conditioned response by strapping the electrodes to the dog's leg so that it could not escape or avoid the shock by flexing its leg. Would this have made any difference? Certainly it would have influenced the frequency with which the leg response was part of a conditioned response. A rather well-known experiment with guinea pigs illustrates this point. In this experiment (Brogden, Lipman, and Culler, 1938) a buzzer was sounded a number of times, each time just before the guinea pigs received an electric shock. These animals were shocked in a revolving cage, and half of them

avoided the shock if they began to run when the buzzer was sounded. The remaining animals were shocked whether or not they ran. The results of this experiment can be seen in Figure 25. As this figure clearly shows, the animals that escaped or avoided the shock learned to run very quickly. Those conditioned exclusively by the classical technique, however, did not increase their frequency of running to the sound of the buzzer very much. This does not mean that they did not give "emotional" conditioned re-

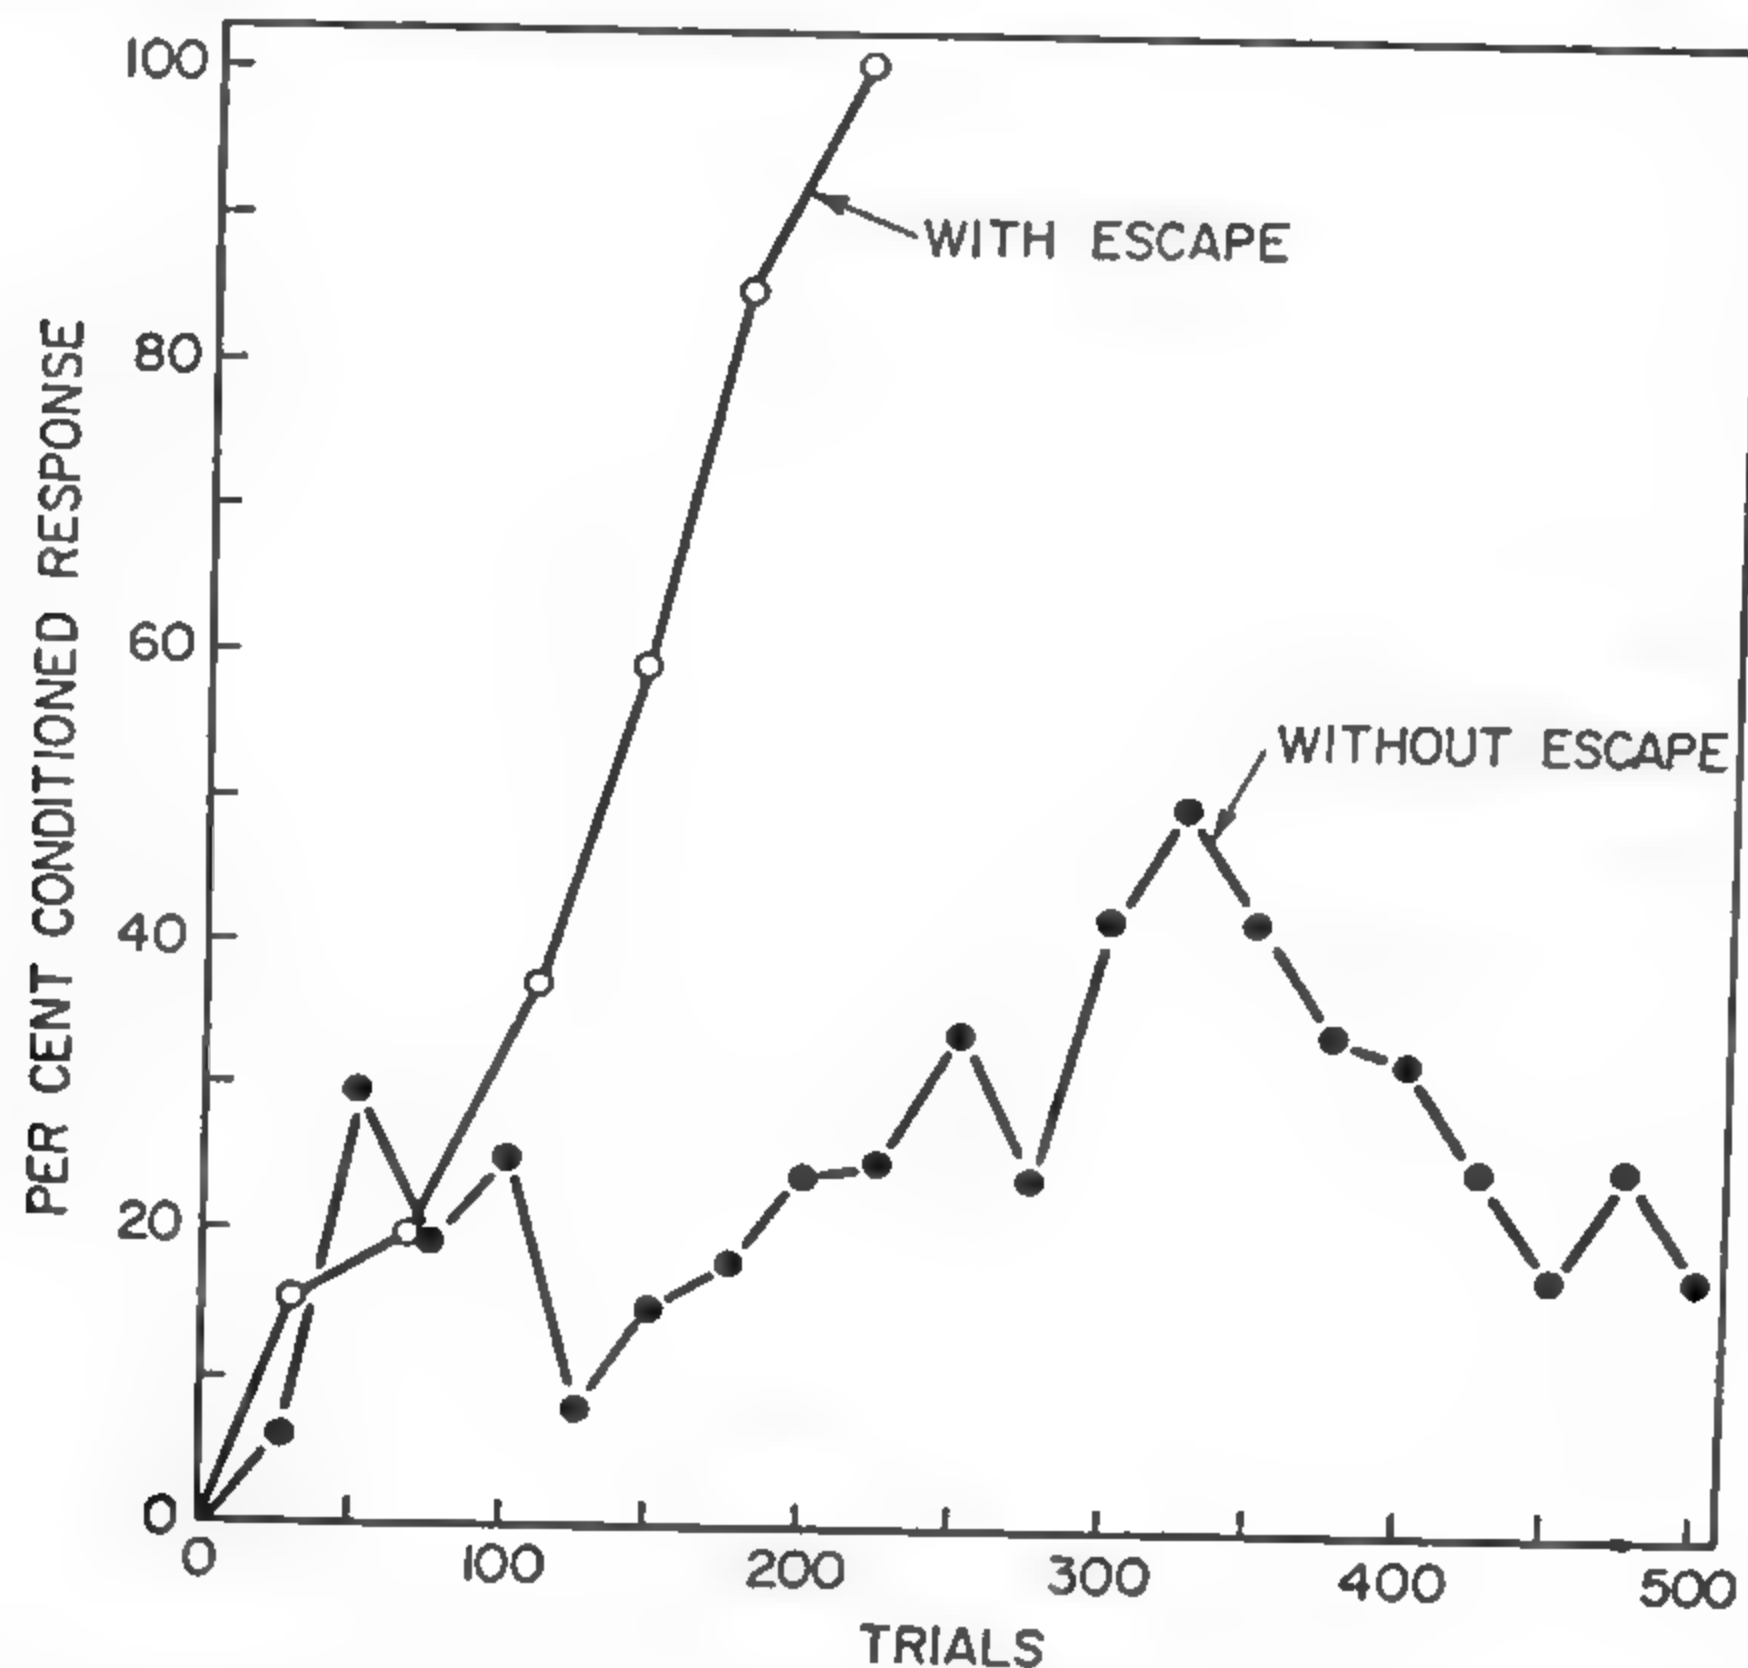


Figure 25. A comparison of classical and instrumental conditioning of an operant. The solid line represents the frequency of conditioned responses to a buzzer when the animals (guinea pigs) are shocked for not running. The broken line represents frequency of conditioned responses when the buzzer is always followed by shock, whether or not the animals respond. (From Brogden, Lipman, and Culler, 1938.)

sponses on these trials, it means merely that the instrumental act of running, since it was not reinforced, did not dramatically increase in frequency. Probably the increase in running that does occur under the classical conditioning technique is the result of the high level of activity that is part of the emotional conditioned response.¹

The paradox of avoidance conditioning. For a long time theorists worried about an apparent paradox in avoidance conditioning. When an animal

¹ There is no implication in this experiment that classical conditioning is necessarily poorer than instrumental conditioning of the same response. See Kimble, Mann, and Dufort (1955).

escapes a noxious unconditioned stimulus such as shock by responding, it is rewarding. But, why should avoidance of the noxious stimulus be rewarding? How can a stimulus not experienced by the animal be said to be a source of satisfaction or drive reduction?

Common sense would say, of course, that the animal comes to anticipate the shock and that is why it responds to the buzzer. The learning theorist is not likely to be satisfied with this answer, however, since he will want to know how and why the anticipation develops. A number of theorists have attempted to answer this question, and one of the most widely accepted solutions is by Mowrer. Let us look at his notion about how avoidance conditioning develops.

Mowrer's two-factor theory of avoidance conditioning. Since Mowrer (1947) argued that avoidance conditioning is a two-stage process, his theory is known as a two-factor theory. He says that first of all the animal learns to give a classical conditioned emotional response to the conditioned stimulus. This, of course, is a matter of pure involuntary association; it does not depend on the animal's voluntarily selecting a response because it is drive-reducing. As we saw in the last chapter, however, emotional disturbance or fear produced by a neutral stimulus is itself a drive. If the animal performs an instrumental response that reduces the emotional disturbance, this is a reinforcing state of affairs and in subsequent situations the animal will again choose this response. Thus, if flexing the leg to escape shock also serves to reduce slightly the emotional responsiveness to the tone, it will be a response preferred by the animal.

In a nutshell then, Mowrer's theory says that the animal avoids the shock not just in order to avoid it, but to escape from the emotional state now conditioned to the buzzer. Actually, Mowrer was able to support his argument with some important evidence. What would happen, for example, if we altered the avoidance situation slightly? Suppose, instead of the buzzer always being sounded until the shock came on, the buzzer was terminated if the animal anticipated the shock by giving an instrumental response. Mowrer and Lamoreaux (1942) showed that, if this were the case, learning of the avoidance response would be more rapid. Thus, if it is possible for the animal to *escape* the conditioned stimulus, learning to avoid the unconditioned stimulus is easier.

Because the buzzer is contiguous with the shock, to summarize the matter, it comes to elicit emotional reactions that are unpleasant to the animal. Some response terminates the buzzer, and the emotional disturbance is partially relieved. This relief constitutes need reduction, which is reinforcing

for the instrumental response. So really, one might say, the animal responds because the response terminates the unpleasant emotional activity conditioned to the buzzer. Such behavior has the appearance of *anticipatory* behavior—and Mowrer believes that he has solved in part the problem of anticipatory behavior—which has always been a stumbling block to learning theory.

We need not necessarily accept the distinction between contiguity conditioning of emotional (primarily autonomic) responses and need-reduction conditioning of instrumental responses, as does Mowrer. The need-reducing properties of the reinforcement for the instrumental response can be thought of simply as the instigator. Mowrer's theory does necessarily imply, however, that the conditioned emotional response must occur before the animal can learn to avoid the unconditioned stimulus by making an instrumental response.

More about the two-factor theory. The two-factor theory has had enormous implications for the theory of behavior. It has suggested, to some bold theorists, that many of the mechanisms of neurotic behavior are learned instrumental responses used by the individual to reduce in part some learned fear or anxiety. Thus it is important that we examine this notion in some detail.

Solomon and Wynne (1954) have extended the two-factor theory to take into account some experimental data they believe to be important. One of these is the problem of extinction of the emotional reaction. If an avoidance habit is well established, the animal no longer receives painful stimulation, therefore the emotional reaction to the conditioned stimulus ought to extinguish and thus indirectly weaken the instrumental response. Solomon points out, however, that things do not always happen this way. Sometimes avoidance responses are extraordinarily resistant to the effects of extinction, and furthermore latencies of the conditioned avoidance response keep getting smaller even after the animal has received its last shock (Solomon, Kamin, and Wynne, 1953). One explanation for this, suggest Solomon and Wynne, lies in the fact that the animals perform the instrumental response extremely quickly when the conditioned stimulus appears; so quickly, as a matter of fact, that there is no time for the conditioned stimulus to elicit the full-blown emotional response. Thus, they argue, the emotional response is *conserved* by extremely rapid instrumental responding. Since this responding does not allow anxiety to develop, it is not fully reinforced, and the animal may lose interest in responding rapidly. Every so often, therefore, the reaction time for the instrumental response is rela-

tively long. It is protracted enough to elicit a full-blown emotional response, and when this happens the animal does perform the instrumental response and is reinforced. In addition, Solomon and Wynne argue that intense emotional conditioned responses are irreversible (not subject to extinction). This is an interesting suggestion, and we shall return to it later.

Some of the further experimental implications of the two-factor theory are being intensively explored. Additional work by Mowrer and his associates (Mowrer and Solomon, 1954; Mowrer and Aiken, 1954) has established the importance of temporal relationships between conditioned and unconditioned stimuli in the establishment of fear, and the freedom of such a response from any drive-reducing mechanism. Some difficulties with the two-factor theory do crop up, however. For example, Dykman, Gantt, and Whitehorn (1956) find that a component of the conditioned emotional response, conditioned heart rate changes, varies rather nicely with the intensity of the painful unconditioned stimulus, but the differentiation of this emotional response *does not* precede the differentiation of an instrumental avoidance response to the conditioned stimulus. Thus, the differences of instrumental responding to different intensities of painful stimulation do not seem to be dependent upon prior differences in conditioned emotional states to these intensities as perhaps they should be according to the theory. We may expect to see a good deal of experimental work on this and related problems in the two-factor theory in the next few years.

THE ROLE OF EMOTION IN CONFLICT

There is a widely held notion that conflict between the anxiety-producing aspects of certain goals and their positive attractiveness is responsible for certain disturbances in behavior. Stated more generally, conflict is regarded as important in behavioral disorders. Consequently there have been rather intensive laboratory studies of the nature of conflict in an effort to establish a relationship between the theory of conflict and more general theories of learning. In this section our aim will be to see how this experimental work has managed to determine a relationship.

The Basic Principles of Conflict

Lewin (1931, 1935) was responsible for the basic analysis of conflict behavior. His analysis, of course, was within the framework of general notions about fields of psychological forces responsible for the psychological state of the organism. Lewin has made a distinction between three types

of conflict. These are (1) approach-avoidance conflict, (2) approach-approach conflict, and (3) avoidance-avoidance conflict. Approach-avoidance conflict occurs when the same stimulus object is both desirable and unpleasant. An example of such conflict could be seen by shocking thirsty rats for drinking water. In approach-approach conflict, the organism is faced with the necessity of making a choice between two equally desirable goal objects—I can spend my savings for a new car or for a vacation. Finally, in avoidance-avoidance conflict, the organism is faced with stimuli that are only to be avoided—I can go to the dentist and suffer in the chair, or I can avoid the dentist and endure a toothache.

N. E. Miller (1944) has examined the implications of these definitions by making certain hypotheses about the relative strengths of approach and avoidance gradients. Miller points out that there are four principles fundamental to the analysis of conflict behavior. These are:

1. The tendency to approach a goal is stronger the nearer the organism is to the goal.
2. The tendency to avoid an unpleasant stimulus is stronger the nearer the subject is to the stimulus.
3. The tendency to avoid increases more rapidly with nearness to the unpleasant stimulus than does the tendency to approach with nearness to a goal.
4. These *gradients* of approach and avoidance are dependent upon the strength of the particular drive. With stronger drives, the tendencies to approach goals or avoid unpleasant stimuli are much stronger.

Brown (1948) was able to show that these principles do apply to the behavior of rats. He trained a group of rats to run down a straight alley to food associated with the presence of a light. The animals wore a light harness which could be attached to a calibrated spring, so that if the animals were momentarily restrained, the strength of pull could be measured. In the experiments, Brown measured the strength of pull at two different points in the alley, at one point near the goal and at one point far from the goal. For a study of the avoidance gradient, the animals were shocked in the goal box and then the strength of the tendency to avoid the goal box was measured at two different points in the alley.

The results of Brown's studies can be seen in Figure 26. The avoidance gradient is steeper than the approach gradient, and increasing the drive, either by making the animals hungrier or by increasing the shock, raises the level of both the approach and avoidance gradients.

Approach-avoidance conflict. If the same goal were made both attractive and unpleasant, one would expect from the results of Brown's experiments that the animals would approach part way and then stop and vacillate because the avoidance and approach gradients had crossed. Where this happened they would be about equal in strength, and thus the animals would be subject to an almost equal pull in opposite directions. This is

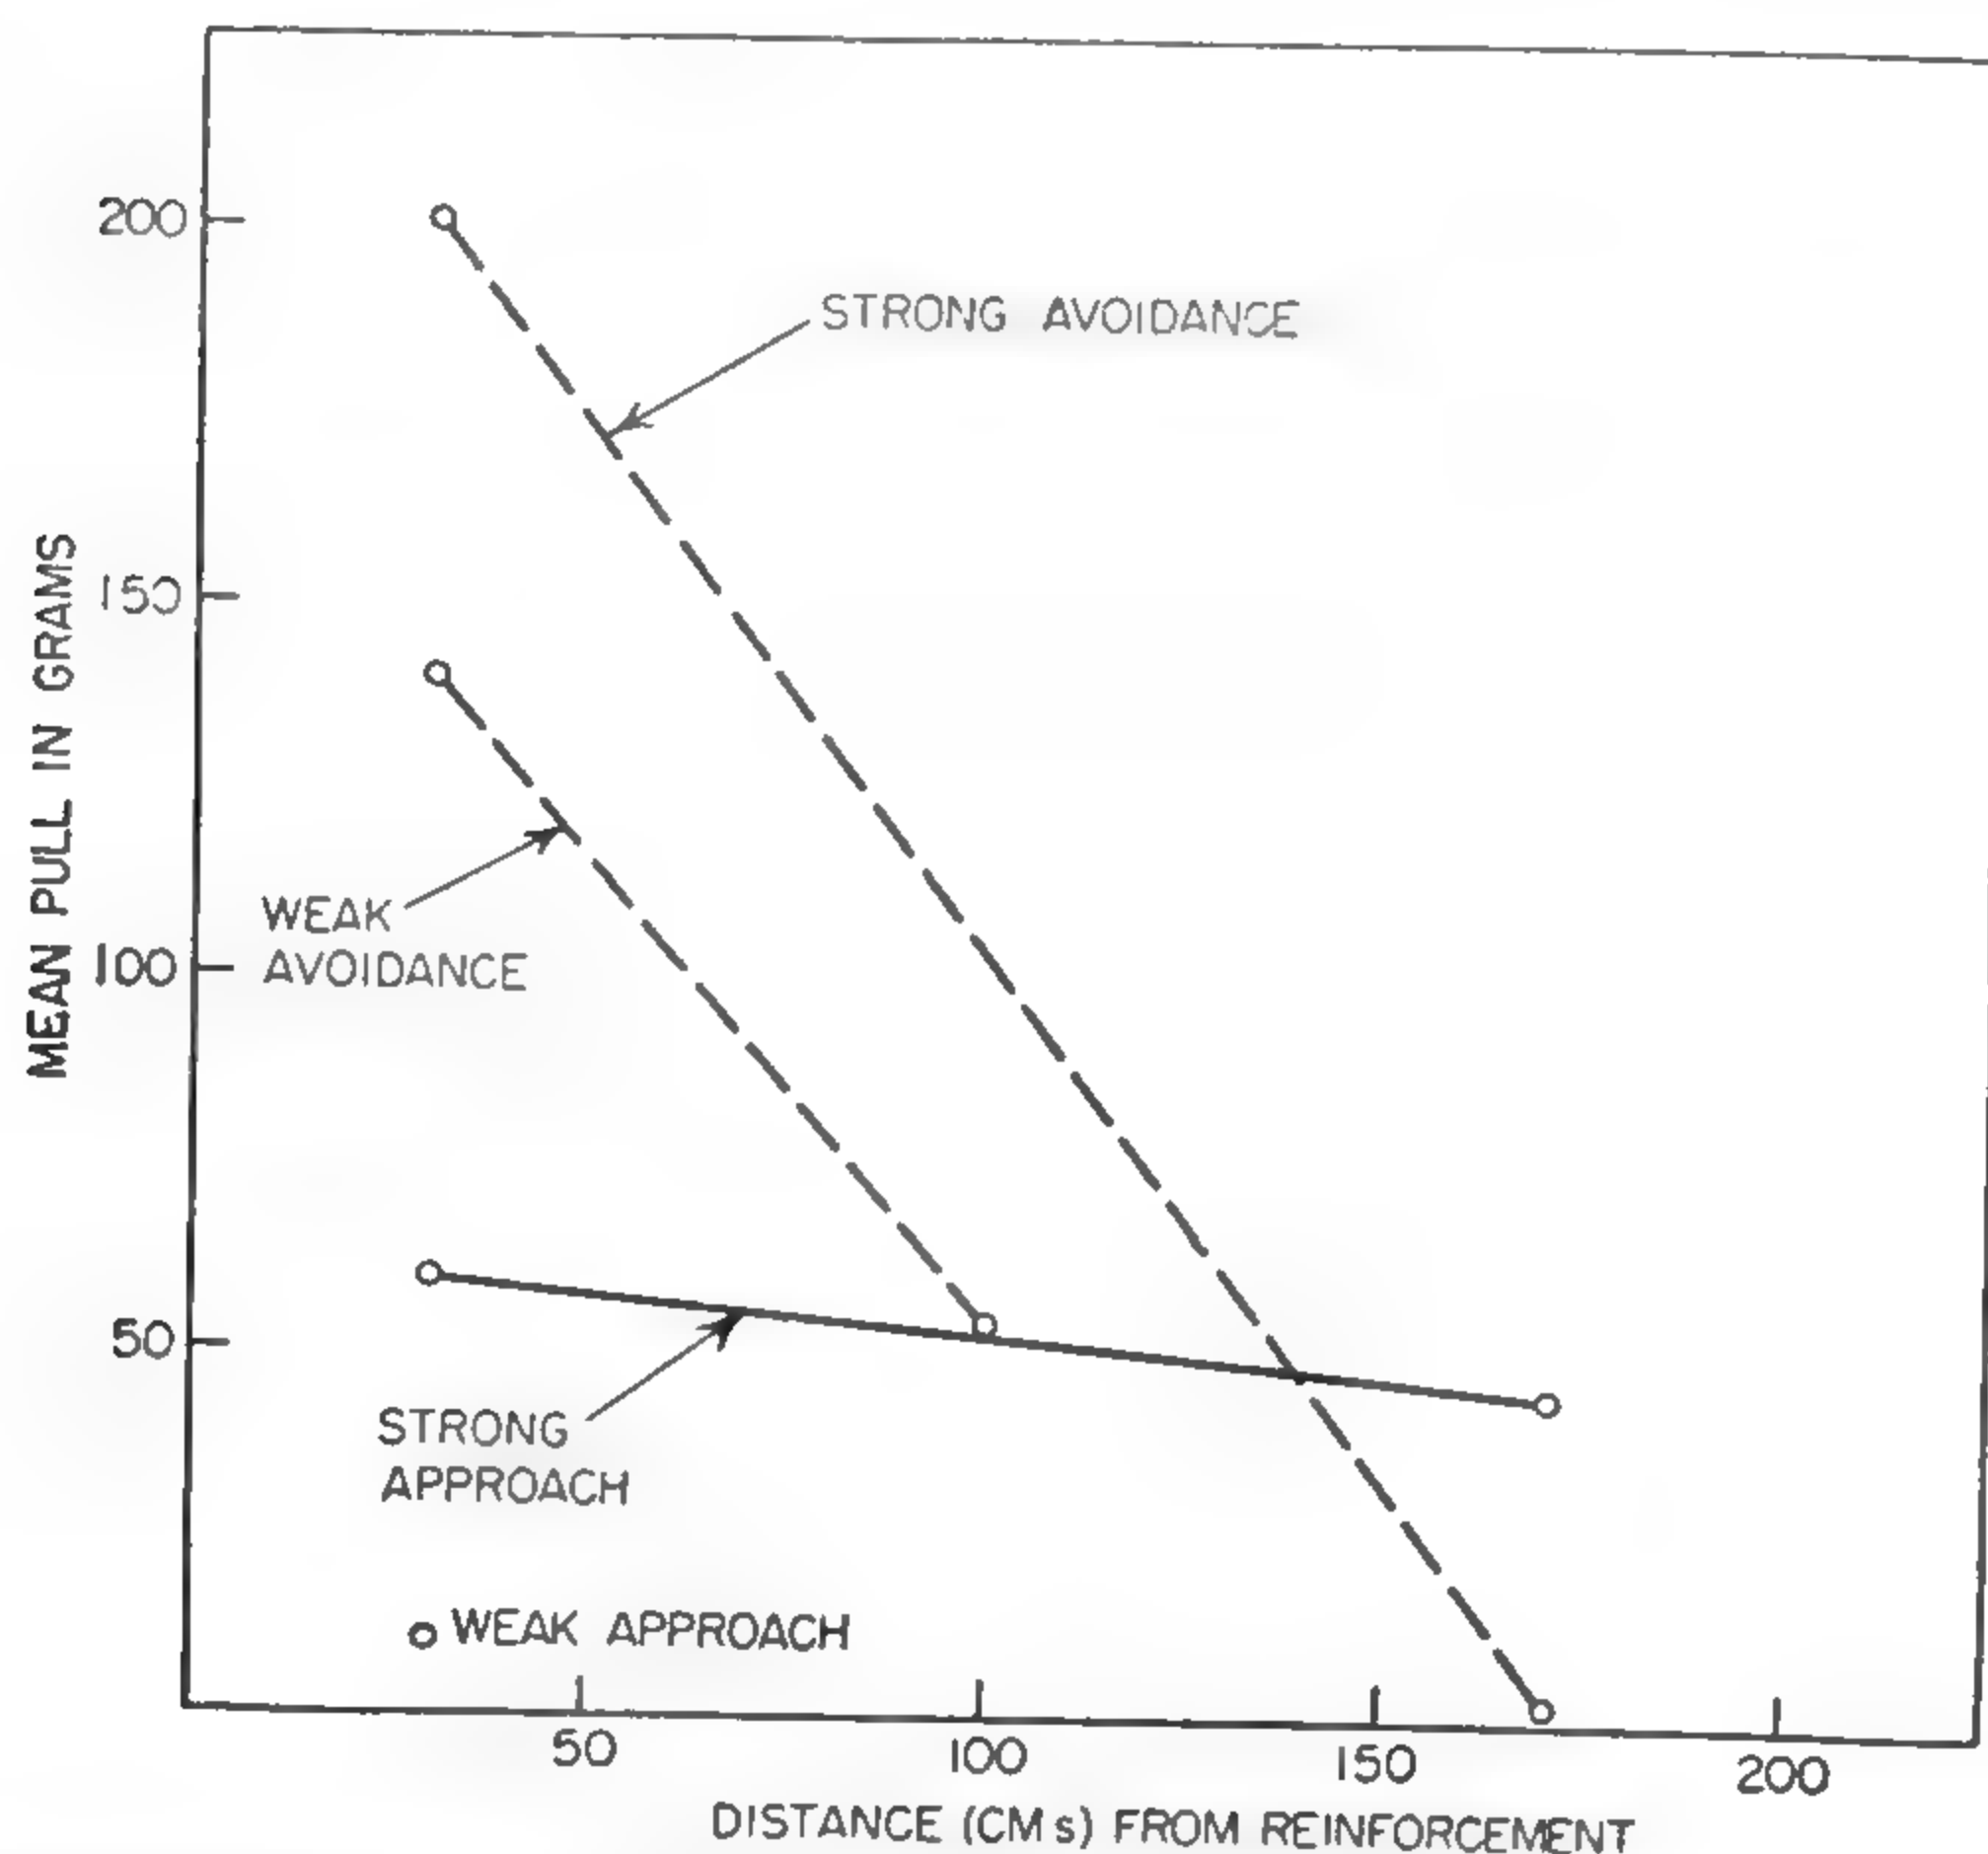


Figure 26. The solid line is an approach gradient. The dashed lines are avoidance gradients of different strengths. Conflict is strongest where the avoidance gradients cross the approach gradients. Note that the strong avoidance gradient crosses the approach gradient further from the goal than does the weak avoidance gradient. (From Miller, 1944; data from J. S. Brown.)

actually what happened in an experiment (Miller, Brown, and Lipofsky, 1943). Animals ran to a goal box in order to get fed, and they were shocked while they ate. Thus, both an approach gradient and an avoidance gradient were produced. The results of the experiment showed that if the shock were very strong, the animals would run only a little way down the alley and then hesitate; if the shock were weak, they would run almost the full length of the alley before they stopped. Thus, the principles tested by Brown are verified in this experiment.

Avoidance-avoidance conflict. The avoidance-avoidance conflict situa-

tion is one of the most interesting of the cases. In this kind of conflict vacillation is maintained the longest. Avoidance-avoidance conflict makes the organism choose between two undesirable alternatives. If it moves away from one of the alternatives toward the second, the avoidance tendency to the second becomes greater and the organism will tend to move back toward the first again. This kind of situation Miller calls a *stable equilibrium*. Two experiments (Klebenoff, 1939; Hunt, 1943) demonstrate clearly the tendency toward vacillation as well as a tendency to try to leave the field.

Approach-approach conflict. Miller points out that approach-approach conflict will produce little vacillation—it is essentially an *unstable equilibrium*. Once the conflict is resolved (perhaps by a purely chance choice), the organism tends to approach one of the goals with increasing strength, since, of course, as it gets closer to that goal the attractive pull of that goal becomes greater and the attractive pull of the alternative goal becomes weaker.

Not all experimental results agree in showing that approach-approach conflict works this way, however. Godbeer (1940) reports considerable vacillation in approach-approach conflict in children. Williams (1943) says that this is because the child can anticipate symbolically the frustration of losing one of the goals, something a rat cannot do. Williams points out that many supposed cases of pure approach-approach conflict also involve an element of avoidance. It is not to be expected that approach-approach situations will produce no vacillation.

Miller recognizes this, and he suggests that approach-approach situations sometimes involve another kind of conflict, a conflict he names “double approach-avoidance.” This occurs because as one goal is approached the other becomes more distant and the organism wishes to avoid giving up this goal. The net result is that the approach tendency may become relatively weak as the organism approaches the goal. The distant goal seems more attractive, and the organism may turn back to it.

Further Applications of the Principles of Conflict

The problem of displacement. One application of the principles of conflict has been to *displacement*. Displacement occurs when an act of aggression cannot safely be directed toward its instigator, but instead gets diverted to some innocent object. Thus the man who is angry at his boss inhibits expressing aggression to the boss and instead, expresses it to his secretary. Displacement would occur only if we assume that the response of inhibiting the aggression (avoidance) shows a much steeper curve of stimulus general-

ization than the displaced (generalized) aggression response itself. Miller and his associates (Miller and Kraeling, 1952; Murray and Miller, 1952) were able to establish that this is the case for a simple maze habit in rats. Thus, there is considerable justification for extending the analysis of conflict to the problem of displaced reactions.

Conflict and anxiety. In conflict that involves avoidance tendencies, there is nearly always some level of conditioned emotional behavior present. This is because, as we saw earlier, the avoidance tendency itself depends upon the conditioned emotional reaction. Usually avoidance means the emotion of fear or of anxiety, so in any conflict involving avoidance one of these is present. In addition there has been a widespread belief that the conflict per se is the source of additional fear or anxiety. We shall examine this notion when we come to look at the problem of experimental neuroses. It serves to point out, however, the great importance of the analysis of the conditions of conflict.

The spatial analogy in the analysis of conflict. Most of the experiments on conflict we have discussed are experiments in which actual spatial distance from the goal was an experimental variable. The space, in these cases, is physical. However, from Lewin has come the notion that conflict can be conceived of in the framework of forces within a *psychological space*. This notion of psychological space is implicit in much of our thinking about conflict. For example, when we speak of avoiding the dentist versus avoiding the toothaches as an example of conflict, we implicitly conceptualize this situation in an imaginary psychological space. We see the individual as a point in this space surrounded by barriers and buffeted by two opposing forces. The analogy to experiments on rats in physical space would have much greater force if we could deal with this imaginary psychological space in the same way that we deal with physical space. This is one of the great unsolved problems in the analysis of conflict, and the fruitfulness of this approach to more general cases of psychological conflict will depend upon the degree of success we have in solving this problem.

THE PROBLEM OF PUNISHMENT

Behavior is punished when some painful stimulation or the threat of such stimulation is made contingent upon it. Thus, if I punish my dog by cuffing it for chewing up the newspapers, in effect I have made the cuffing contingent upon chewing newspapers. Usually punishment is administered to moderately well-motivated behavior in order to stop the behavior. Thus

punishment is a technique for one organism controlling the behavior of another, though, of course, some particular behavior of an animal or human being may be inadvertently punished by a physical accident. In this case a pseudocorrelation between the behavior and its punishing consequences is set into the organism.

The essential property of punishment, then, is that it is painful or unpleasant stimulation applied to moderately well-motivated or highly motivated behavior. The rationale behind the social use of punishment is that it will stop undesirable behavior, though our modern insights into personality mechanisms lead us to suspect that punishment is sometimes administered to satisfy some need of those who inflict it. As a problem in the theory of learning, however, punishment is best seen as the influence of painful or unpleasant emotion-producing stimulation upon behavior.

The Classical Experiments on Punishment

For many years experimental and theoretical work on punishment was dominated by the views of the late E. L. Thorndike. He was, for most of his professional life, an educational psychologist, and consequently his concern with the problem of punishment was influenced by the formal demands of the classroom. In a word, he was concerned with the influence of verbal praise and reproof upon behavior. Since his views and experiments were influential, let us look briefly at them.

Thorndike's views and experiments. In his earliest writings, Thorndike simply adopted the common-sense view that punishment would reduce the tendency to repeat behavior it immediately followed (weakened connections, as Thorndike put it). Later, however, Thorndike (1932a) examined the problem more thoroughly and came to another conclusion. He decided that punishment really had no inherently weakening effect on behavior. Let us see how Thorndike arrived at this rather surprising conclusion.

Thorndike performed a number of experiments on punishment, with both animal and human subjects. As his experiments on human subjects are better known, let us review one of these. Thorndike (1932b) asked subjects who were not familiar with Spanish to guess the correct English equivalent to a Spanish word. He gave a series of such words in the form of a multiple-choice test. If the subject chose the correct word, the experimenter informed the subject by saying "right"; if the word was incorrect, the experimenter said "wrong." It was possible to find the influence of "right" or "wrong" by finding out, on further testing, if the subjects gave the pun-

ished or rewarded responses with a frequency greater or lesser than that expected by chance. Since there were five alternatives for each item, the frequency of repetition by chance would be 20 per cent. This was the base line against which Thorndike measured the effect of saying "right" or "wrong." As it turned out, "right" did increase the tendency to repeat the initially given response. The consequence "wrong", however, rather than being weakening, seemed to Thorndike to have a slight strengthening effect.

Criticisms of Thorndike. Many similar experiments by Thorndike (1935) led him to the same conclusion. Other investigators leveled serious criticisms at him, however. First of all, they argued, the use of "chance" as a base line against which to measure the effects of reward or punishment is not good; it is extremely unlikely that responses given by subjects in such tests are at the chance level; rather responses go in "runs," or sequences, to which the one-in-five chance expectancy would not apply. Stephens (1934), for example, used an empirically determined probability of repetition by including items to which the experimenter gave no reward or punishment. The frequency of repetition without any consequences turned out to be 36 per cent instead of the 20 per cent assumed by Thorndike. Measured from this base, punishment did have a slight weakening effect (given the mildness of the punishment and the low motivation for repetition we should not expect a large effect).

Punishment indirectly weakens. Thorndike argued, however, that wherever punishment appeared to weaken a response it was an indirect effect. This happens because punishment produces variability of behavior. In maze learning, for example, punishment of an error may lead the learner accidentally to make the correct response, and this response indirectly eliminates the punished response, because the correct response is strengthened due to its effect. It has never been precisely clear what Thorndike meant by the indirect weakening of punished responses (see the controversy, Postman, 1947; Stone, 1948) but it would be something like this: Punishment may or may not weaken a response, but it clearly cannot be the mirror image of the action of a reward. If a response is rewarded, it is apparent to the organism in a typical experimental test that repetition of *this* response will again be rewarding, but if a response is punished, it is not clear to the organism which of the other available responses will be rewarded. Thus, behavior after punishment will be more variable, e.g., the prediction of the exact response to be made by the organism will be more difficult. This is the most general statement we can make. In the Thorndikian situations the

punishment is a mild one, and perhaps its main function is to provide information. If this is so, reproof for making a particular response does not supply as much information as praise, for it does not tell the subject which response is correct.

Punishment in Experiments on Animal Behavior

Many of the important conditions of punishment are ignored in the Thorndikian experiments. The magnitude and quality of unpleasant effects are ignored, and, perhaps more importantly, the effects of restricting the available responses are ignored. This last point needs some amplification.

In Thorndike's experiments the learner may satisfy his motivations by making responses other than the ones punished. He is seldom interested in defending his original response as the "right" one, and he will readily abandon it when he knows it is "wrong." Thus, the effect of punishment is to produce variability of behavior, as Thorndike suggested.

What happens, however, when the only response which satisfies a drive is punished? The child who is punished for taking cookies from the cupboard knows no other way to satisfy a craving for sweets. No matter how many raps on the knuckles the child may receive, the cookies lose none of their attraction, and he is not likely to find another way of getting cookies. There are interesting experiments upon the behavior of rats that bear directly on this problem.

Skinner's experiment. Skinner (1938) studied the effect of punishment upon the lever-pressing response in the Skinner box. In his experiment, the punishment was a sharp slap on the paws from the lever itself. To measure the effects of punishment, Skinner obtained extinction curves following periodic reinforcement with food from four rats. Of these rats, two were slapped at the beginning of the extinction period and two were not. The result was that the rats slapped at the outset of extinction showed a much lower rate of responding during extinction. As a matter of fact, the punishment seemed to inhibit the response entirely for a period of time. The curious thing, however, is that by the time extinction had been completed, the slapped rats had caught up, in total number of responses emitted, with the rats not slapped. These results can be clearly seen in Figure 27.

Skinner interpreted the results of this experiment to mean that punishment has only a temporary effect upon behavior. It does not affect the "reserve" of responses which the animal will eventually emit; it affects only the rate at which this reserve will be exhausted. Early in extinction the punished animals show a lower rate of responding, but eventually the rate for these

animals accelerates so that they show a total number of responses equal to that of the unpunished animals.

Estes' experiments. Estes (1944) followed Skinner's experiment with an intensive investigation of the factors affecting punishment of an isolated instrumental response—lever pressing in the rat. In these experiments he substituted an electric shock for the slap. The shock was delivered through the lever, so that rats used to receiving food as the result of lever pressing, sometimes received shock instead.

Estes' first conclusion is much like Skinner's. The effect of punishment is primarily upon the rate of responding rather than on the over-all tendency to emit a particular response. Estes did find, however, that if the punishment were intense enough or lasted long enough, there was a slight permanent depression in the reserve of available responses. Nevertheless punishment was never able to eliminate the extinction curve completely—

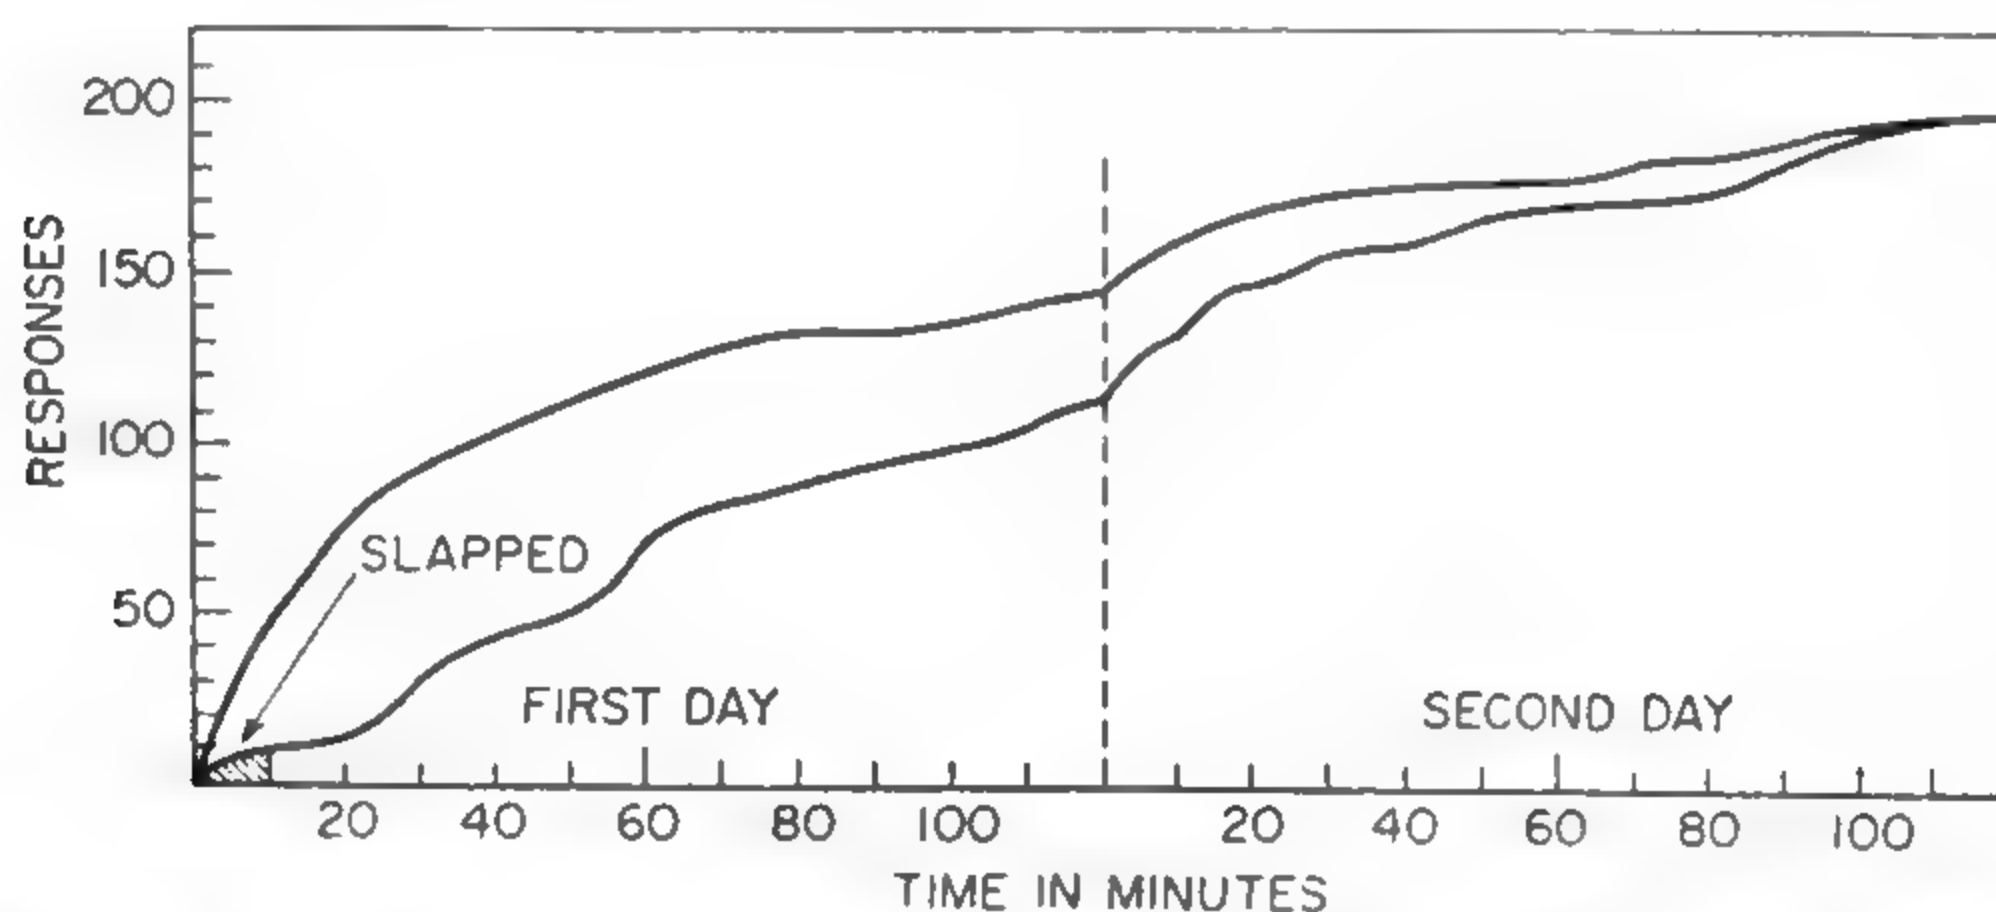


Figure 27. The effects of punishment upon resistance to extinction. The initial effect of punishment is to reduce the rate of responding. There is, however, recovery, and eventually the total number of responses is about the same for the punished and control animals. (Skinner, 1938.)

rats always came back to respond after punishment, even though they were never again reinforced. Interestingly enough, when Estes punished his rats periodically instead of for every response, the depression of the rate of responding was not nearly so severe, though it lasted longer. Thus, the results of periodic, or partial, punishment are somewhat analogous to the results of partial reinforcement.

Punishment and conditioned emotional responses. Estes interprets the effects of punishment somewhat as follows: anxiety or fear aroused by the shock becomes conditioned to the stimuli in the experimental box. This anxiety, most strongly conditioned to the visual stimulus of the lever, causes the animal to withdraw. Withdrawal reduces the anxiety and is therefore

reinforcing. In Estes' terms, anxiety becomes prepotent over the lever-pressing response, which consequently disappears for a time. However, since the animal does not press the lever, the anxiety conditioned to the lever extinguishes and the animal gradually returns to pressing the bar. Unless the response is again punished, the anxiety gradually diminishes, allowing the animal to exhaust the reserve of lever-pressing responses it built up during food reinforcement.

Thus, in this experiment, the effectiveness of punishment seems to come from an emotional reaction conditioned to the environment in which the rewarded instrumental response is likely to occur. The emotional reaction prevents, or at least reduces, the tendency for the emission of the instrumental response. When the emotional reaction extinguishes, as it should if the animal is not punished again, the tendency to perform the instrumental response returns.

Punishment, then, seems to be a rather special case of avoidance learning. This is pointed out by Dinsmoor (1954) and he shows that a fruitful analysis of the possible effects of punishment can be obtained by studying experiments in avoidance learning. He brings out that the operation of punishment is essentially a two-step process as it is in avoidance learning.²

The use of the suppression period produced by punishment. It is clear that punishment may not be too effective in completely eliminating some behavior. Punishment may momentarily weaken a response, but it does nothing to eliminate the behavior permanently. This conclusion seems to be in accord with the clinical and anthropological evidence about aggressive behavior punished by parents or by some agent of society. The tendency toward aggression is not removed until it can be brought to free expression, and thus special psychotherapeutic techniques are frequently devised for the express purpose of allowing aggression to occur in a displaced (therapeutic) situation.

It is possible, however, to use the temporary period of response inhibition produced by punishment to teach an organism a new method of reaching the same goal. An experiment by Whiting and Mowrer (1943) suggests that this may be a valuable technique in the control of behavior.

These investigators studied ways of getting rats to abandon old, preferred

² Hunt and Brady (1955) have compared the effectiveness of punishment and a conditioned emotional response in suppressing instrumental behavior. Their results show that the conditioned emotional response is more effective, because of the more adequate stimulus control it exerts over behavior.

adjustments in favor of new ones. They used three groups of rats. All three were taught to run a simple maze. After the habit of running to the goal box had been acquired, they extinguished it in one group of rats by omitting the reinforcement. They continued to reinforce the habit for the second group, but placed a barrier in the path leading to food. In the case of the third group, they punished the running response by an electric shock. The situation was then changed so that all three groups could find the way to the goal by an alternative route. The animals that had been punished on the original habit showed the fewest errors in reaching the goal by the alternative pathway. They also showed less regression to the original habit after the alternative habit had been learned.

Thus, teaching an alternative response during the period of suppression induced by punishment could be an effective means of redirecting behavior. Of course, as we might expect, if the second habit is extinguished, there will be a tendency to revert to the original, punished habit.

The effects of traumatic punishment. When we were looking at the problem of avoidance learning, we saw that Solomon and Wynne (1954) suggested that sometimes the conditioned emotional reaction set up by painful stimulation was extremely resistant to extinction—indeed perhaps never could be extinguished. If the same thing applied to punishment, it is possible that a punished response would never recover its strength. The experiments on avoidance learning that Solomon and Wynne were talking about were experiments in which very traumatic levels of electric shock had been used. In the experiments on punishment we have looked at thus far, the punishment was relatively mild. Consequently, it is possible that under a really traumatic punishment, the instrumental response punished may never recover its original strength.

This seems to be the case in some studies of Masserman (1943) on punishment in cats. As we shall see shortly, the experiments of Masserman suggest that the effects of such punishment may go far beyond the instrumental response punished; there may be changes in almost all behavioral adjustments of the animals. This is the case in the experimental production of neurosis, the problem we shall consider next.

EXPERIMENTAL NEUROSIS

One of the interesting problems in the relationship between emotion and learning is provided by those investigators who have attempted to pro-

duce disturbances in behavior by experimental conditions in the laboratory. In this section we shall look at some of the experiments and their resultant theories.

The Experimental Production of Neurosis in Animals

Pavlov's experiments. Some observations by Pavlov (1927; 1928) provide the introduction to the technique of producing "experimental neurosis." While studying the conditioned salivary response, workers in Pavlov's laboratory noticed that animals sometimes developed long-standing disturbances of behavior. Pavlov began to study these disturbances seriously after an experiment in which he attempted to condition salivation by using electric shock as a conditioned stimulus. The electric shock seemed to work fairly well for a while, but after an attempt at generalization, the conditioned salivation was replaced by a violent struggle quite foreign to the usual behavior of the dogs.

A subsequent experiment led Pavlov to an intensive study of disturbances of behavior. This experiment started out as an attempt to test the limits of form discrimination in a dog. The animal was trained to salivate to the presentation of a luminous circle on a screen. Then an attempt was made to get the animal to discriminate between a circle which was reinforced and an ellipse which was not. After discrimination was established, the ellipse was gradually changed to look more like the circle. For a while, the dog continued to discriminate, but then as the two stimuli became more alike, discrimination suddenly became very poor. In addition, the whole character of the animal's behavior changed. The dog, which had been trained to stand quietly in the apparatus, began to struggle and squeal. Eventually it became quite violent and resisted being taken to the experimental room. This behavior is what Pavlov called an experimentally produced neurosis.

Pavlov believed that the primary cause of this disturbance of behavior in the test situation was a conflict in the cerebral cortex between opposed forces of excitation and inhibition. The excitation was produced by the positive conditioning of one stimulus and the inhibition by the extinction of the other. Both conditioning and extinction occur in discrimination, and when the stimuli get close to each other, the conflict between the opposing forces of excitation and inhibition becomes strong enough to disturb the whole cortex.

Pavlov's general notions about cerebral physiology have more or less fallen into disrepute as we have come to know more about the facts of cerebral activity. His interpretation of neurosis in terms of conflict, how-

ever, is still influential. If we were to rephrase Pavlov's theory in language less tainted with outdated physiology we might say that experimental neurosis occurs because of the inability of the animal to tell when to respond and when not to. This conflict could be the result of a discrimination that is too difficult for the animal. It should be noted in passing, however, that such a conflict can arise from causes other than difficult discrimination. If the animal must delay a conditioned response too long after the presentation of the conditioned stimulus (because of a delay in presentation of the unconditioned stimulus) there will be a conflict between the tendency to respond and the tendency to inhibit the response until the appropriate time.

Other investigations of experimental neurosis. Since Pavlov's original observations there have been many other attempts to produce disturbances in the behavior of animals in the laboratory. Gantt (1936) has closely followed Pavlov's work, and he has been able to confirm Pavlov's observations on disturbances of behavior during difficult discrimination.

In studies by Liddell (1944) the conditioned flexion response has been used rather than the salivary response. In experiments on sheep, dogs, and pigs, Liddell and his associates have been able to produce many behavioral disturbances in the laboratory as well as disturbances in bodily functioning. In these experiments there has been great emphasis placed upon the role of restraint of the experimental animals during testing. Restraint, monotony of stimuli, and other factors are responsible for the disturbances, according to Liddell. It is interesting to note that Liddell finds that many of these disturbances persist when the animals are removed from the experimental situation.

Some of the most interesting observations of experimentally produced disturbances of behavior come from Masserman (1943). Masserman used an instrumental response in producing his disturbances. Cats were trained to open a small food box after a stimulus was presented. The techniques used to produce the disturbances of behavior were frustration, conflict, and punishment. Suffice it to say that by these techniques Masserman was able to produce markedly deviant behavior in his cats. Some cats became "phobic" toward the box, particularly after punishment. Other cats went into semicataleptic states, refusing to move or even to eat for long periods of time. Masserman attempted a variety of therapeutic devices to relieve the neurosis after it had been established. Rest was about the most effective means of allaying the animals' symptoms. Also effective were reducing the hunger drive through forced feeding outside of the experimental room, "reassurance," and social imitation of nonneurotic cats.

Theories of Experimental Neurosis

Today there is less emphasis upon trying to find analogies between disturbed behavior in animals and that in human beings than there was a few years ago. Nearly all psychologists realize that a technique used to produce a disturbance in an animal need not achieve the same effect in human beings. There is less concentration upon techniques which produce disturbed behavior in animals than upon the production of conflict between two habits or upon the production of an inhibiting fear in animals. Let us look at some of the notions about disturbed behavior in animals, not with the idea of finding analogies to similar behavior in human beings, but with the idea that we might be seeing simple models of some of the processes at work in human neuroses.

Masserman's theory. Of all of the recent theories of behavioral disturbances based upon experimental studies of animals, Masserman's comes closest to considering disturbances in animals to be analogous to those in human beings. He emphasizes the fact that traumatic situations in the experimental tests are responsible for phobic reactions and fixations, much as they are supposed to be in the case of human personalities. The symptoms of withdrawal and confusion exhibited by his animals, Masserman believes, occur because of the conflict between the motive to get food and the anxiety associated with the punishment of the food-getting responses.

While Masserman does not explicitly point up the relationship, his analysis comes very close to the kind we should expect from our earlier discussion of basic learning mechanisms in emotional behavior. In addition, however, he points out the many parallel points between disturbed personal and social adjustments in human beings and what he believes to be their counterpart in animal behavior. Though Masserman's views of the mechanisms in disturbed behavior do not provide us with a systematic theory, they are important in that they represent the first real attempt to bring together psychiatric and psychoanalytic theories with the experimental study of animal behavior.

The Dollard and Miller theory. A systematic account of the relationship between problems of emotional learning in animals and behavior disorders in human beings is presented by Dollard and Miller in *Personality and Psychotherapy* (1950). These writers are not specifically interested in presenting a theory of experimental neurosis, but they are interested in applying the basic principles of learning to the description of human personality disorders. Though their aim is somewhat different, the general re-

sult of bringing together experimental observations of animal behavior and observations of human beings in trouble is the same.

Perhaps the most important principle that Dollard and Miller use is approach-avoidance conflict. People frequently find themselves in the unfortunate position of having conditioned fear responses aroused by situations which can reduce their positive drive. In other words, there is conflict between fear and desires. Dollard and Miller discuss in detail the case history of an individual who suffered from conflict between the fears instilled by a repressed childhood and her normal, adult sexual desires. Since the fear actually prevented this woman from satisfactorily satisfying her positive needs, misery and anxiety arose. Furthermore the fear produced repression of fear-arousing thoughts and images so that the fear itself was prevented from undergoing extinction. The net result was a vicious circle in which no opportunity was given for the reduction of fear or the elimination of conflict. This, say Dollard and Miller, is the general scheme for disturbances in personality and the production of psychosomatic symptoms.

Dollard and Miller also apply the principles of learning to the process of readjustment of maladjusted people. They present an account of the course of psychotherapy in one individual and point out how each aspect of psychotherapy works in terms of the principles of learning. They show, for example, how talking about a problem in a permissive and relaxed atmosphere could result in extinction of the fear responses and consequent resolution of a conflict. They likewise interpret the transference that takes place during therapy as a special case of stimulus generalization, and show how discrimination learning must take place if the individual is to distinguish between appropriate and inappropriate behavior in various life situations.

Maier's theory. The position advanced by Dollard and Miller really needs no new principles to describe the mechanisms in disturbed behavior; the learning principles discussed thus far in this book are perfectly adequate. We should, however, look briefly at one theory of experimental neurosis that does invoke a new principle. This is the theory presented by Maier (1949). In experiments by Maier and his associates, rats are required to discriminate between two visual stimuli presented on the Lashley jumping stand. The stimuli are reinforced and punished in a random manner, so that the rats really cannot learn or find a solution to the problem. Animals frequently refuse to choose in this situation, even though they may be very hungry. Thus it is sometimes necessary to "drive" the animals to choose by stimulating them with a blast of air (which emits a piercing sound painful to rats) or an electric shock. When this happens and the rats are forced to

respond, they usually fixate on a particular response—jumping to the left card or jumping squarely at the middle, missing both cards. Once this response becomes fixed, it never changes. This Maier has called “abnormal fixation.”

The animals in these experiments, Maier says, are frustrated, and this is responsible for the fixation. Animals, he finds, that have developed a fixation may even refuse to jump to an open window in which food is displayed. Rather, they prefer their “fixated” response. Maier believes, moreover, that his experiments show that behavior elicited during a state of frustration has certain unique properties. Frustrated behavior is not motivated behavior in the usual sense; it solves no problems and has no goals. Aggression, regression, and fixation are all characteristics of frustrated behavior. Since behavioral disturbances can be described in terms of aggression, regression, and fixation, Maier believes that a major portion of the disturbances in behavior reported in studies on experimental neurosis are due to frustration.

The strength of the fixations reported in experiments on the Lashley jumping stand is remarkable. Making available the solutions to the problem does not break up the fixation (Maier, 1949), and even training the rats to walk to the correct stimuli does not eliminate it (Feldman, 1953). Strong punishment, persistently administered, however, does seem to break the position stereotypes of fixations (Knopfmacher, 1953).

Thus Maier introduces a new concept, frustration induced fixation, and he, as well as others, has produced a good deal of evidence to support this notion. It is an interesting one, and only time and new evidence will tell us whether such a concept is necessary or whether all the things it purports to explain can be clarified by the ordinary principles of learning. Perhaps the greatest importance of the notion, so far as we are concerned at the moment, is that it suggests that the laws of adaptive behavior break down under the stress of frustration.

Some comments on experimental neurosis. We can see, from this brief survey, that many basic principles in the psychology of learning have been put to work answering questions about the causes of disturbed behavior. We have also seen that experimental work on animals sometimes leads us to formulate new principles which, in turn, may be applied to problems of disordered behavior in human beings.

It should be pointed out that there is a vast gulf between the disturbed behavior of animals in the laboratory and the disturbed behavior of individuals in trouble. In most studies of experimental neurosis, disruption of behavior has been produced by stimuli that elicit strong emotional re-

sponses. These disturbances, as well as the emotional behavior, are situational. The animals may be disturbed in the laboratory but they are less or not at all disturbed when out of the laboratory. Human neuroses are not situational. Thus, in the matter of behavioral disturbances the rat has the advantage over the human being. The rat does not carry its anxieties around; rather its anxieties are conditioned to specific stimuli, and it shows signs of anxiety only in the presence of these or similar stimuli. The human being—through the mediation of verbal behavior—carries his anxieties about and uses implicit verbal behavior to reinforce them. Thus, perhaps the major stumbling block to a direct analogy between disturbed behavior in laboratory animals and in human neurotics is the low capacity animals have for engaging in symbolic and linguistic behavior. Human beings can much more readily generalize and extend their fears and conflicts because they can symbolically face fear of punishment or conflict even when the objective situation does not demand it. This fact and others make us aware of the enormous gulf that exists between experimental neuroses of the laboratory and disturbed personalities. We should not ignore, however, the extent to which animal behavior in the experimental laboratory can provide us with models of basic mechanisms in personality maladjustments. The extraordinary degree to which the results of simple experiments can be used to describe personality processes is well illustrated in Dollard and Miller's previously mentioned *Personality and Psychotherapy* (1950). Nor have Dollard and Miller stood alone in this endeavor; Mowrer (1950) and others have explored much the same course, and our understanding of basic mechanisms in the dynamics of individual personality has been enriched by their work.

STRESS AND LEARNING

Before we leave the topic of emotion we have one important problem to consider, the effects of emotional activity upon learning and the performance of learned skills. Emotional activity arises when the organism is under some stress. Examples of stress are familiar enough—the navigator of a modern bomber might be troubled by the threat of physical injury and the need for quick, precise performance of his task to ensure the success of his mission.

Emotion and Motivation in Stress

It is obvious that stress involves components other than emotional activity. It is a state in which some motive or motives remain ungratified; this

may mean failure to avoid or escape threat (fear) or it may even mean lack of satisfaction of a powerful appetitive desire (sex, for example). Thus stress means heightened motivation as well as heightened emotional activity. What is more, in experimental studies of human behavior, it is practically impossible to produce emotional disturbance without also producing a change in motivation (Lazarus, Deese, and Osler, 1952). Thus, we cannot examine the effects of emotional arousal upon learning and skilled performance without taking into consideration the effects of change in motivation. This point was raised in the last chapter in the section on experimental studies of human motivation. A number of studies were reviewed in which investigators had tried to relate individual differences in motivation level to learning and performance on various tasks. The difficulty with these studies, remember, was lack of agreement as to whether or not the measures of individual differences used by these investigators actually assessed motivation. Some workers point out that these studies really were concerned with the relationship between performance and individual differences in emotionality and performance.

Unfortunately there is no way out of this dilemma at present. Perhaps, from a strictly operational point of view, we have no justification for talking of the effects of emotion and motivation as being independent of each other, at least as far as human skilled performance is concerned. There is evidence from other sources, however, that emotional activity and motive states do derive from somewhat independent sources, and given the present general viewpoint in psychology, we find it convenient to talk about emotion and motivation separately.

Roughly stated, level of motivation is responsible for level of performance. In other words, higher motivation leads to more adequate performance if the latter leads to eventual satisfaction of the motive. Emotional arousal provides activity which for all but a small number of skilled acts is likely to produce interference. Trembling, gastrointestinal upset, etc., are all things that may distract from or directly interfere with some particular task an individual is trying to perform. Consequently, in the vast majority of specific instances, we should suppose that emotional arousal would have a detrimental effect upon skilled performance.

What, then, of the effects of stress? Stress, as we saw, involves both motive and emotional components. One very possibly would bring about improved performance, and the other, deteriorated performance. What actually happens when individuals are stressed in an experimental test and the effects of this situation upon performance evaluated?

Experimental Studies of Stress and Performance

As the confusing roles of emotion and motivation in stress would lead us to imagine, there is no simple way of stating unequivocally that stress will have specific effects upon learning and the performance of particular skilled acts. There are few stable generalizations that can be made.

One which seems to receive support from a wide variety of experimental investigation is that stress increases the speed with which repetitive tasks are carried out (Lindsley, 1946; McKinney, et al., 1951). This increase in speed is generally accompanied by a higher rate of errors. The net result is that some people actually improve their performance because of the greater speed, and others decrease in performance because of an increase in errors (Lazarus and Eriksen, 1952).

Furthermore, we suspect that individual reactions to stress are determined in part by stable personality patterns. In the last chapter we discussed some studies of the relationship between performance on skilled tasks and scores on personality inventories. We saw that these personality inventories probably assess individual differences in motivational as well as emotional factors. Therefore, a thoroughgoing exploration of the relationship between personality inventory scores and behavior under stress ought to provide us with some interesting and analytic data about stress and personality. A bare beginning to this study has been made (Deese, Lazarus, and Keenan, 1953; Osler, 1954). Perhaps studies of personality characteristics and behavior under stress will provide a means of separating the motivational and emotional components in the reaction to stress.

There is tentative evidence that poorly organized habits are more susceptible to deterioration from stress than well-established habits. Information items on an intelligence test, for example, are little influenced by stress, but the immediate memory span (memory for things heard or seen just once) is susceptible to stress (Moldawsky and Moldawsky, 1952). Also, it seems likely that performance early in the learning of a new skill is influenced more by stress than is later performance (Deese and Lazarus, 1952).

Beyond these simple statements it is difficult to go. It is likely that a more thorough analysis of the effects of stress upon learning and performance will depend upon new experimental techniques and new ways of assessing individual differences in emotional reactivity and motivational level.

CHAPTER 7

SERIAL LEARNING

A basic fact of behavior is that it is sequential; acts follow one another in a continuous stream. One of the central problems for the psychology of learning is to consider how this stream of behavior becomes organized. As we try to solve this problem we shall draw upon many basic principles already discussed, as well as upon new principles and ways of studying behavior. Therefore, this chapter is about the possible applications of the basic principles of instrumental learning to particular facts of serial or sequential patterns of behavior, and the new techniques and methods used in describing such patterns.

Many of the examples of learning we looked at in the earlier chapters involved sequences of responses rather than single responses. We treated these examples, however, as if they were single responses. In the Skinner-box experiments, for example, the learned behavior was characterized as a "lever-pressing" response, and we treated it as if only one act were necessary to carry out the learned instrumental response. Actually, there is an elaborate chain of responses involved in pressing the lever, a fact pointed out by Skinner (1938). The rat must see the lever, and this may lead to lifting the paws; it cannot begin the actual depression movement until it feels the lever; it must then seize and eat the food. Skinner described at least four separate responses as part of the chain involved in pressing the lever, and it is probable that a really detailed analysis would yield more. Certainly parts of the chain change as learning progresses (Hurwitz, 1954).

We treated the instrumental activity in the Skinner box, however, as a single response, since we were not interested in the parts of the chain of behavior and how they became organized. The "response" was defined as something which produced an effect upon the environment—moving the

lever through a certain excursion. After so defining the response, it was possible to study the effects of different variables and to treat these as if they were affecting only one response rather than a whole chain of responses. Before we can even begin to appreciate the full range of problems in learning, however, we must examine the available evidence on the organization of response sequences—how responses interact with one another when they occur in the stream of behavior.

In one sense, the larger portion of the remainder of this book is devoted to the problem of the organization of response sequences, since interactions between responses provide the basis for much of forgetting, problem solving, and transfer of training. Therefore, as we consider each of these problems separately, we shall also be examining certain special effects of response sequences.

SERIAL LEARNING IN ANIMALS

Let us first look at the problems of serial learning in animals. Since animal behavior is less complicated than human behavior and since animals have little available to them in the way of symbolic processes or linguistic behavior, we shall find applying the basic principles of behavior to serial learning in animals relatively simple. Later, when we come to deal with human serial learning, we shall examine the new methods and principles.

Techniques in Studying Serial Learning in Animals

Everybody knows that psychologists spend a good deal of time teaching white rats how to run mazes. This is partly because the maze provides a good technique for the comparative study of animal behavior, but also because the maze is an important source of information about serial learning in animals. There are, of course, other techniques for the study of serial processes in animals, but since the maze is the most widely used and important technique, let us look at it first.

The maze. The maze is the tool par excellence for the study of the behavior of rats. The rat's natural habitat is a mazelike series of burrows, and it learns mazes efficiently and well. Consequently, most of the studies of maze learning have used rats as subjects.

There has been an enormous variety of mazes used in experimental studies. Figure 28 shows the floor plan for one maze. The rat is placed in the box marked S, and it must make its way to the reinforcement in the box marked G. At each of the choice points, the rat must choose between a

blind alley and the path that will eventually bring it to the goal. An error is counted when a rat enters one of the blind alleys. In this particular maze the correct sequence of turns is a relatively simple LRLR. The maze can be lengthened or shortened by additional choice points, and instead of having these arranged for systematic alternation, they can be arranged randomly.

Other mazes can be more haphazardly made, so that, for example, the distance between successive choice points is not the same. The famous

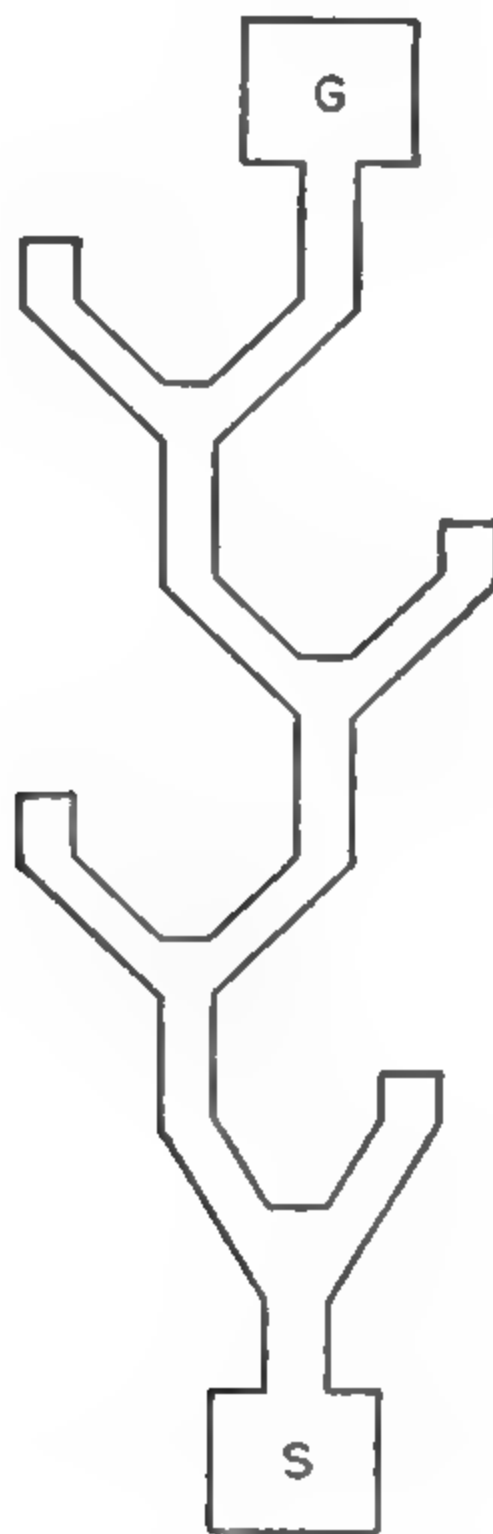


Figure 28. The floor plan of a multiple-choice-point maze. The number of choices is probably smaller than would be used in practice. Each entry into a blind alley would contribute one error to a trial.

Hampton Court maze is of this sort. Still other mazes will have more than one blind alley at a choice point. Finally, some have alleys leading away from the choice point at different angles; a right angle is typical (as in a multiple T maze), but more acute angles are frequently used.

The usual animal maze is an enclosed alley. In some experiments, however, the animals are allowed to run on open platforms. In enclosed mazes, of course, animals cannot see more than one choice point ahead, but in an elevated one, it is possible for the animal to obtain some overview of the maze pattern.

There is obviously a wide variety of mazes, each one suitable for a different problem. One simple maze, the linear maze, has been important in

theoretical studies of learning. Figure 29 shows an example of a linear maze. It is just a long alley which has been divided into three or four compartments. There are anywhere from two to four doors between each compartment. Only one of the doors will swing open; the others will be locked shut. The rat must learn which door is free to open and to run to it. An important feature of the linear maze is that the choice points closely resemble one another; consequently, stimulus generalization from one choice point to another is likely to be much the same.

Other techniques. Instead of presenting a maze in which the animal must run, it is perfectly feasible to provide a sequence of other kinds of instrumental activity. One experimenter (Arnold, 1947), for example, presented rats with a series of buttons which had to be pushed in an invariant order; the measure of learning in this case was taken from reaction times, or latencies.

Frequently a Skinner box is modified so that the rat is presented with

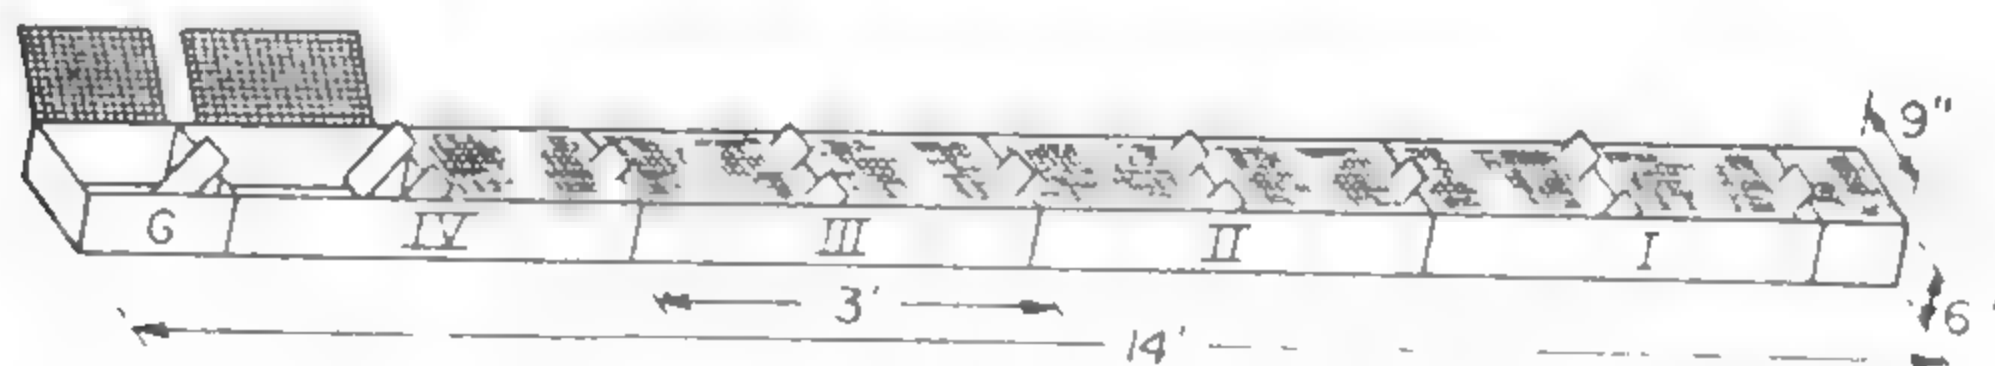


Figure 29. A linear maze. There are four choice points with two choices at each. There is a "centering" door between each choice point to prevent the favoring of one door over the other by virtue of the previous correct response.

two levers instead of one. In some cases the rat must press the two levers in some particular order—two presses on the right lever and then two presses on the left lever (Schlosberg and Katz, 1943). In other cases the object is just to see how the rat distributes its choices between two levers. In this connection, some experimenters have studied how the rat distributes its choices between pressing the lever and going to the food tray by putting the food tray on the opposite side of the box from the lever (Frick, 1953).

The particular techniques for studying serial learning will vary with the problem being studied and the organism being tested. The maze is the instrument of choice with the white rat, because this animal is capable of learning an extended series of reactions in the maze. With other organisms, other techniques are sometimes preferred, and if one wishes to study the natural sequences of behavior, the maze has limited usefulness since it forces the animal to perform its choice reactions in an invariant sequence.

Let us first consider the events that occur when an animal is forced to learn in the unchanging pattern imposed by a maze.

Theories of Maze Learning

The original investigations of maze learning (Small, 1899, 1900) were made with the idea that the rat does not learn a series of movements but, rather, locations in the maze. This common-sense notion that the rat learns to recognize places has always been held by some investigators of animal behavior, but, for a time at least, the predominant view seemed to be that rats learned sequences of movements. Today, we think it rather unimportant whether the rat learns sequences of movements or locations in the maze, and we suspect that actually both kinds of learning occur (Blodgett and McCutchan, 1947). Historically, however, many important ideas were allied with the notion that rats learn sequences of responses in the maze, so let us see how this hypothesis arose.

Watson's theory of maze learning. Perhaps the earliest attempt to account for the facts of maze learning in a systematic way is that of J. B. Watson. Watson was one of the great pioneers in the study of learning and one of the most important of the early behavior theorists.

Watson assumed that an animal learning to run a maze is learning a sequence of movements. After Pavlov's experiments on the conditioned salivary response became generally available in this country, Watson assumed that each of the movements in the sequence was much like a Pavlovian conditioned response. (Notice that he was applying the principles of classical conditioning to instrumental learning.) Thus, Watson said, the basic process involved in habit was that by which a neutral stimulus comes to elicit a response when paired with a stimulus which already has the power to elicit the response. Watson's position in this matter was so important and fundamental that we can do no better than look at a quotation from the second edition of his famous textbook (Watson, 1924, pages 293-294).

Certainly at birth or shortly thereafter the elements or units out of which every habit is formed can be noted. We mention the contraction and flexion of the fingers, or of the lower and upper arm, raising and lowering of the hand, rotation of the head, bending the trunk from side to side, the backward and forward, well-systematized movements of the legs, and a host of others. The conclusion is forced upon us that in habit no new elementary movements are needed. There are enough present at birth and more than will ever be combined into complex unitary acts. . . . One needs only to examine the five or six day old infant to be reasonably convinced that there is no need for the formation of additional reflex arcs to account for all later organization. The new learned element in habit is the tying together or integration of separate movements in such

a way as to produce a new unitary activity. And by unitary activity we mean nothing more than the everyday acts of life, such as reaching out the hand for an object that stimulates the eye, picking the object up and carrying it to the mouth or laying it on the table. . . .

We can define habit then as we did instinct as a complex system of reflexes which functions in a serial order when the child or adult is confronted by the appropriate stimulus.

Watson applied these general notions about the development of complex acts to the analysis of maze learning in rats—an example of how a whole series of discrete movements becomes organized into a unitary act. In a complicated maze, Watson thought, the rat must learn a series of responses; for example: run 5 feet, turn left, run 3 feet, turn left again, etc. He thought of the animal which had learned the maze correctly as having synthesized a whole chain of conditioned responses. The ingenious feature of Watson's notion was that the conditioned stimulus for each phase of the sequence of acts came from the preceding act. The conditioned stimuli in maze learning stemmed primarily from the kinesthetic stimuli of the movements themselves. Thus, the "feel" of each movement was the conditioned cue for the next movement.

Watson (1907) thought he had evidence to show that the main sensations rats used in learning mazes were kinesthetic. Unfortunately, it turned out that Watson's experiments were faulty and that kinesthetic cues are not necessarily the most important ones for rats in maze learning (see Munn, 1950). Over and above this fundamental error, there were several other critical features against Watson's notion of maze learning as a simple chaining of conditioned responses.

If Watson's basic idea is correct (aside from whether the conditioned stimuli are kinesthetic or not), the whole chain of responses ought to be very mechanical; each response should be exactly the same each time and follow exactly in its proper sequence. Experimental support for this view was offered by Watson (Carr and Watson, 1908). Rats were trained in a maze in which the alleys could be lengthened or shortened without changing the pattern of turns which the animal had to make. Carr and Watson noticed that animals trained on short alleys (a turn required after a relatively short straightway) attempted to turn at the same place when the alley was lengthened. Thus it looked as though the animals had learned a sequence of movements in which the n th movement was a turn. Apparently they had learned to turn right not when the first opening presented itself but after making a certain number of movements.

Such observations supported the notion of kinesthetic chaining as the basis of maze learning. Subsequent experimental work has shown, however, that this cannot be the only, much less the principal, way in which rats learn mazes. The number of experimental refutations of this notion is large, but perhaps the simplest is provided by Macfarlane (1930). He trained rats to run a maze in which the passages contained water to the depth of 8 inches. The rats were thus forced to swim the maze. He then switched them to walking by putting in a false floor slightly below the surface of the water. Since there was no increase in errors when the animals were transferred from swimming to running, the learned maze pattern did not consist entirely of a chain of conditioned movements.

Watson believed that the most important condition determining learning and retention of the correct pathway in mazes was that correct responses were generally more frequent and also more recent than incorrect responses (Watson, 1914). Watson was not able, however, to account for many of the special features of maze learning with these notions of frequency and recency, and the result has been that more modern theories have tended to abandon them in their simple form.

Hull's analysis of maze learning. Hull, like Watson, has placed great emphasis upon the conditioned response. He also indiscriminately applies the same principles to classical and instrumental conditioning. Unlike Watson, however, Hull does not rely exclusively upon movement-produced conditioned stimuli. The result is that Hull has provided us with a theoretical analysis of maze learning that, while defective in certain particulars, has been the most powerful analysis made to date.

Hull (1930) differs from Watson in that he allowed external cues from the maze to be conditioned to each response. The visual stimuli at a choice point, he said, provide the cues for the response of turning to the right. He conceived of this process as being like classical conditioning, but it does not do violence to his notions to describe turning as instrumental activity brought under the control of a particular discriminative stimulus by differential reinforcement. Hull follows Watson to the extent of assuming that kinesthetic cues can provide part of the chain that keeps responses in their proper order. Moreover, since animals are usually motivated during runs through the maze, he said that there is a drive stimulus which is conditioned to every response the animal makes (Hull, 1930, 1931); in other words, since the animal is equally hungry at every point in the maze, the stimuli from drive are conditioned to every part of the maze.

An essential element in maze learning, according to Hull (1932) is the gradient of reinforcement (Chapter 2). Hull calls this the *goal gradient*, and it refers to the gradient of the effectiveness of reinforcement through the maze. In general, it states that correct responses close to the goal will be reinforced sooner (and hence performed more adequately) than correct responses distant from the goal. Hull (1943) largely attributes this goal gradient to a diminishing effect of secondary reinforcement as the animal gets further from the goal. The goal gradient, of course, is a temporal gradient; it functions as a spatial gradient in the maze only because of the invariant order of choice points in the maze.

One thing the concept of the goal gradient allows Hull to do is deduce that the maze will be learned. This is no mean accomplishment; one of the real deficiencies of Watson's notions of frequency and recency is that they did not unequivocally predict that the rat will learn the maze, a fact pointed out by Thorndike (1915). However, it is clear that any time an animal makes an error, the reinforcement of that error will be further removed in time than the reinforcement of the correct response at the same choice point, for if an animal enters a blind alley it takes it some time to get out. The net result is that correct responses will be more strongly reinforced than incorrect ones. Another deduction from the goal gradient is that if the rat is allowed access to two paths to the goal box, eventually it will come to prefer the shorter path.

A third prediction, one that has not always met with success, is that blind alleys near the goal will be eliminated before blind alleys at the beginning of the maze. Spence (1932) has shown that this is generally true, but there are many specific cases in which the backward elimination of errors is obscured by other factors, such as the direction (relative to the goal) in which a particular blind is pointing.

There is one other mechanism postulated by Hull that is important in the analysis of maze learning. This is the *fractional anticipatory goal response*. We have already seen that Hull suggested that a persisting stimulus, probably arising from the drive state, is conditioned to every response in the maze. It is most strongly conditioned to the goal response, that is, to the response of eating food. It is possible, however, that this persisting stimulus may prematurely elicit the goal response at some point earlier in the maze. This fractional anticipatory goal response allows Hull greater freedom in accounting for the data of maze learning. He can, for example, use this notion in deriving the prediction that animals learn different pathways in

the same maze depending upon whether they are hungry or thirsty (1952). Also, this notion leads to the correct prediction that when maze behavior is extinguished, it does not extinguish serially as one might expect, but seems to disintegrate all at once.

The fractional anticipatory goal reaction is much more of an *ad hoc* notion than that of the goal gradient, and sometimes Hull and his students have used it too freely in explaining the results of experiments. Hull uses this concept to describe much activity which other investigators have usually named anticipatory or cognitive.

Studies on the linear maze. Hull and his students have given a good deal of attention to the study of behavior in linear mazes, or more generally, to situations in which a series of identical or nearly identical responses are required. The linear maze, as we described it earlier, is one in which all choice points look exactly alike. These points usually consist of a number of doors, only one of which will open into the next compartment of the maze.

Since the choice points in a linear maze look so much alike, there is a great deal of stimulus generalization from one to another. Because some of the choice points are closer together than others—choice point II is closer to III than to IV—some will generalize more to one than to others. Thus, if the animal learns to go through the second door at choice point I, it will generalize this tendency rather strongly to II, somewhat less to III, and much less to IV. Likewise, if the third door is correct on choice point IV, the animal will strongly generalize the tendency to go to the third door on choice point III; it will generalize this tendency a little less on point II and very little on I.

We can easily see that these tendencies, based upon generalization, lead to two kinds of errors at any given choice point. Anticipatory errors are those which are incorrect at the choice point at which they occur but correct later on. Thus the rat may choose door 1 when it should have chosen door 2, but 1 is correct at the next choice point. In this case, the rat has anticipated the second choice point. Likewise, perseverative errors are responses which are incorrect where they occur but which were correct at some earlier choice point.

The effect of anticipatory and perseverative errors has been studied in the linear maze under a wide variety of conditions. In some experiments the correct door at each choice is different. In other experiments the same door is correct at every choice point. In some experiments the animals were reinforced only at the end of the maze; this procedure allows the familiar goal

gradient to operate. In other experiments, however, the animals were reinforced after each choice in the series; this produces a radically different gradient of reinforcement.

The results of all of these experiments have been compared with the theoretical deductions derived from Hull's general theory of learning. In most cases the experimental results and the theory agree. We can see very simply and intuitively how the results of these experiments could support the theory.

It is easy to see, for example, that when a different door is correct at each choice point, the anticipatory and perseverative tendencies would produce nothing but errors. Moreover, since the middle choice points are influenced by both ends, the greatest number of errors during learning would be in the center of the maze. This is actually the case in experimental results with the linear maze (Hull, 1947). If, when a different door is correct at each choice point, there is reinforcement only at the end of the maze, there are fewer errors at the end of the maze than at the beginning. If each choice is reinforced, however, there is much less difference in errors between the two end ones, and in actual fact, there are fewer errors at the beginning than the end (Hull, 1948).

If the same door is correct at every choice point, anticipation and perseveration are beneficial to learning, for however the animal anticipates or perseverates at any choice point, it still makes the correct choice. Since perseveration cannot occur at the first choice point and anticipation cannot occur at the last one, anticipation and perseveration are greatest in the middle; therefore, there are fewer errors in the middle (Sprow, 1947; Gladstone, 1948).

In addition, we have some experiments (Smith, 1949; Woodbury, 1950) that show cases in which the choice points are neither all the same nor all different. In these experiments, the point at which the maximal number of errors occurred was in accord with the hypotheses of Hull. For example, if the first three choice points are all the same, but the last one different, then the greatest number of errors occurs at the next to the last choice point, since the animal anticipates the last choice point most strongly here.

Remember that anticipation and perseveration are special cases of stimulus generalization. The rat in the linear maze, when faced with a choice point, will not always be able to tell which compartment it is in. This is a failure of discrimination and is therefore stimulus generalization. Hence, these studies on the linear maze tell us that the principles derived from a study of single instrumental and classical conditioned responses are useful

in analyzing a relatively complex form of behavior. These studies show the way in which generalization of stimuli associated with particular responses may interact to produce a complex pattern of behavior.

Other serial problems. The analysis of maze learning presented by Hull and his associates has been generalized to all sorts of problems. Hull (1952) treats in detail sequences of responses in a modified Skinner box and a special problem, *double-alternation*. The important thing about his treatment is that essentially the same rules apply to these cases as apply to maze learning. In other words, it is necessary to consider generalization from one response to the next, gradients of reinforcement, and all the other mechanisms of interaction between responses. The way in which these factors combine to produce interaction between responses, of course, depends upon the particular problem. The cases which Hull has considered have nearly always been simplified so that the interaction between the various factors is not too complicated. For example, in the linear maze all choice points are alike so that the same amount of stimulus generalization goes from one choice point to the next.

There are, of course, many peculiar features in particular serial acts such as maze performance, that cannot be accounted for by mechanisms we have been describing thus far. Many of these are peculiar to the rat or even to the rat in particular mazes. Consequently, they are not of much general interest to the psychology of learning. There is a thorough discussion of factors limited to specific cases of maze learning in Munn's *Handbook of Psychological Research on the Rat* (1950); some of these factors are important to the comparative study of behavior.

Statistical Sequences of Behavior

In the maze, animals must perform responses in a more or less invariant sequence—they must first make a choice at the first choice point, then at the second, etc. In recent years, however, psychologists have begun to study natural behavior sequences in which animals can emit responses in a variety of orders. In doing so, they have examined in greater detail some examples that we had come to think of as invariant sequences of responses. For instance, we have treated behavior in the Skinner box as if it were much more uniform than it really is. We have more or less assumed that each time the rat presses the lever it immediately turns to the food tray for the reinforcement. This is by no means true. Even when the food tray is placed directly under the lever, the rat will occasionally press the lever several times before eating the food. This may even be the case when the pellet

dropped must be eaten before further pellets are released. Interesting problems arise when we deal with these naturally occurring sequences of behavior, and it will be worth our while to examine them in some detail.

Statistical behavioristics. First of all, we need a method for analyzing the sequential interaction of responses. Such a method is provided by what has been called *statistical behavioristics* (Miller and Frick, 1949). The essence of this method involves computing the probability of a particular response occurring after some other response has been emitted. We may ask, for example, what is the probability of turning to the food tray after lever pressing? Obviously, it is high, though not perfect, because sometimes a second lever press will follow the first. What is the probability of returning to the food tray after eating? It is obviously low, for the rat, rather than return to the empty tray, will press the lever again. Thus, it is clear that the probability of turning to the food tray does not stand by itself, *but depends upon what the organism did just before*. Thus, we may speak of the *dependent probability* of a response, and this, in general, means that the probability of a particular response depends upon what preceded that response.

By analyzing the dependent probabilities in sequences of behavior we can discover many things. We can, for example, compute an index of stereotypy, which tells us how regular and repetitive patterns of behavior are. Let us look at an example of the application of conditional, or dependent, probability analysis to a problem in trial-and-error, or multiple-response, behavior.

Dependent probabilities and trial-and-error behavior. To show how an analysis of dependent or conditional probabilities applies to sequences of responses in trial-and-error behavior, Miller and Frick (1949) have compared the behavior of a rat and a child in a multiple-choice experiment. The problem for the rat and the child was to get out of an enclosure with four doors. Since the door which permitted an exit varied from trial to trial, it was useless for the learner to attempt to try again the door which had been correct on the previous trial. The original purpose of this experiment, which is a part of an older study (Hamilton, 1916), was to get an extended sample of trial-and-error behavior in an insoluble problem.

Miller and Frick point out that if one attempts to predict which of four possible doors the child will pick in any one trial, the best guess is chance—unless information is available on the relative frequency of previous responses and on the response immediately preceding the response in question. Once we have this information, however, we can increase our predic-

tion over chance. If, for example, a sequence of two responses in the child is considered, we find that of the 16 possible sequences (four doors times four doors), 6 sequences appear 80 per cent of the time. Thus a knowledge of the choice just before the response greatly increases the accuracy of our prediction.

When Miller and Frick analyzed the data involved in sequences of three responses, they found great regularity in the third choice. It was nearly always a response which was different from the preceding two. Thus, if the child chose doors 1 and 2, there was a very low probability that either doors 1 or 2 would be chosen as the next response. Miller and Frick point out that this is one respect in which the child differs markedly from the rat. The rat was much more likely to go back to one of the doors chosen on the preceding two responses. This suggests a longer memory span for the child than for the rat.

Dependent probabilities in the Skinner box. A few paragraphs back we said that the sequence of lever pressing, then turning to the food tray, and then back to the lever was not inevitable in the Skinner box. Sometimes, animals press the lever two or three times before going to the food tray. A detailed analysis of sequence in the Skinner box has been made (Frick and Miller, 1951) by altering the box as shown in Figure 30. The food

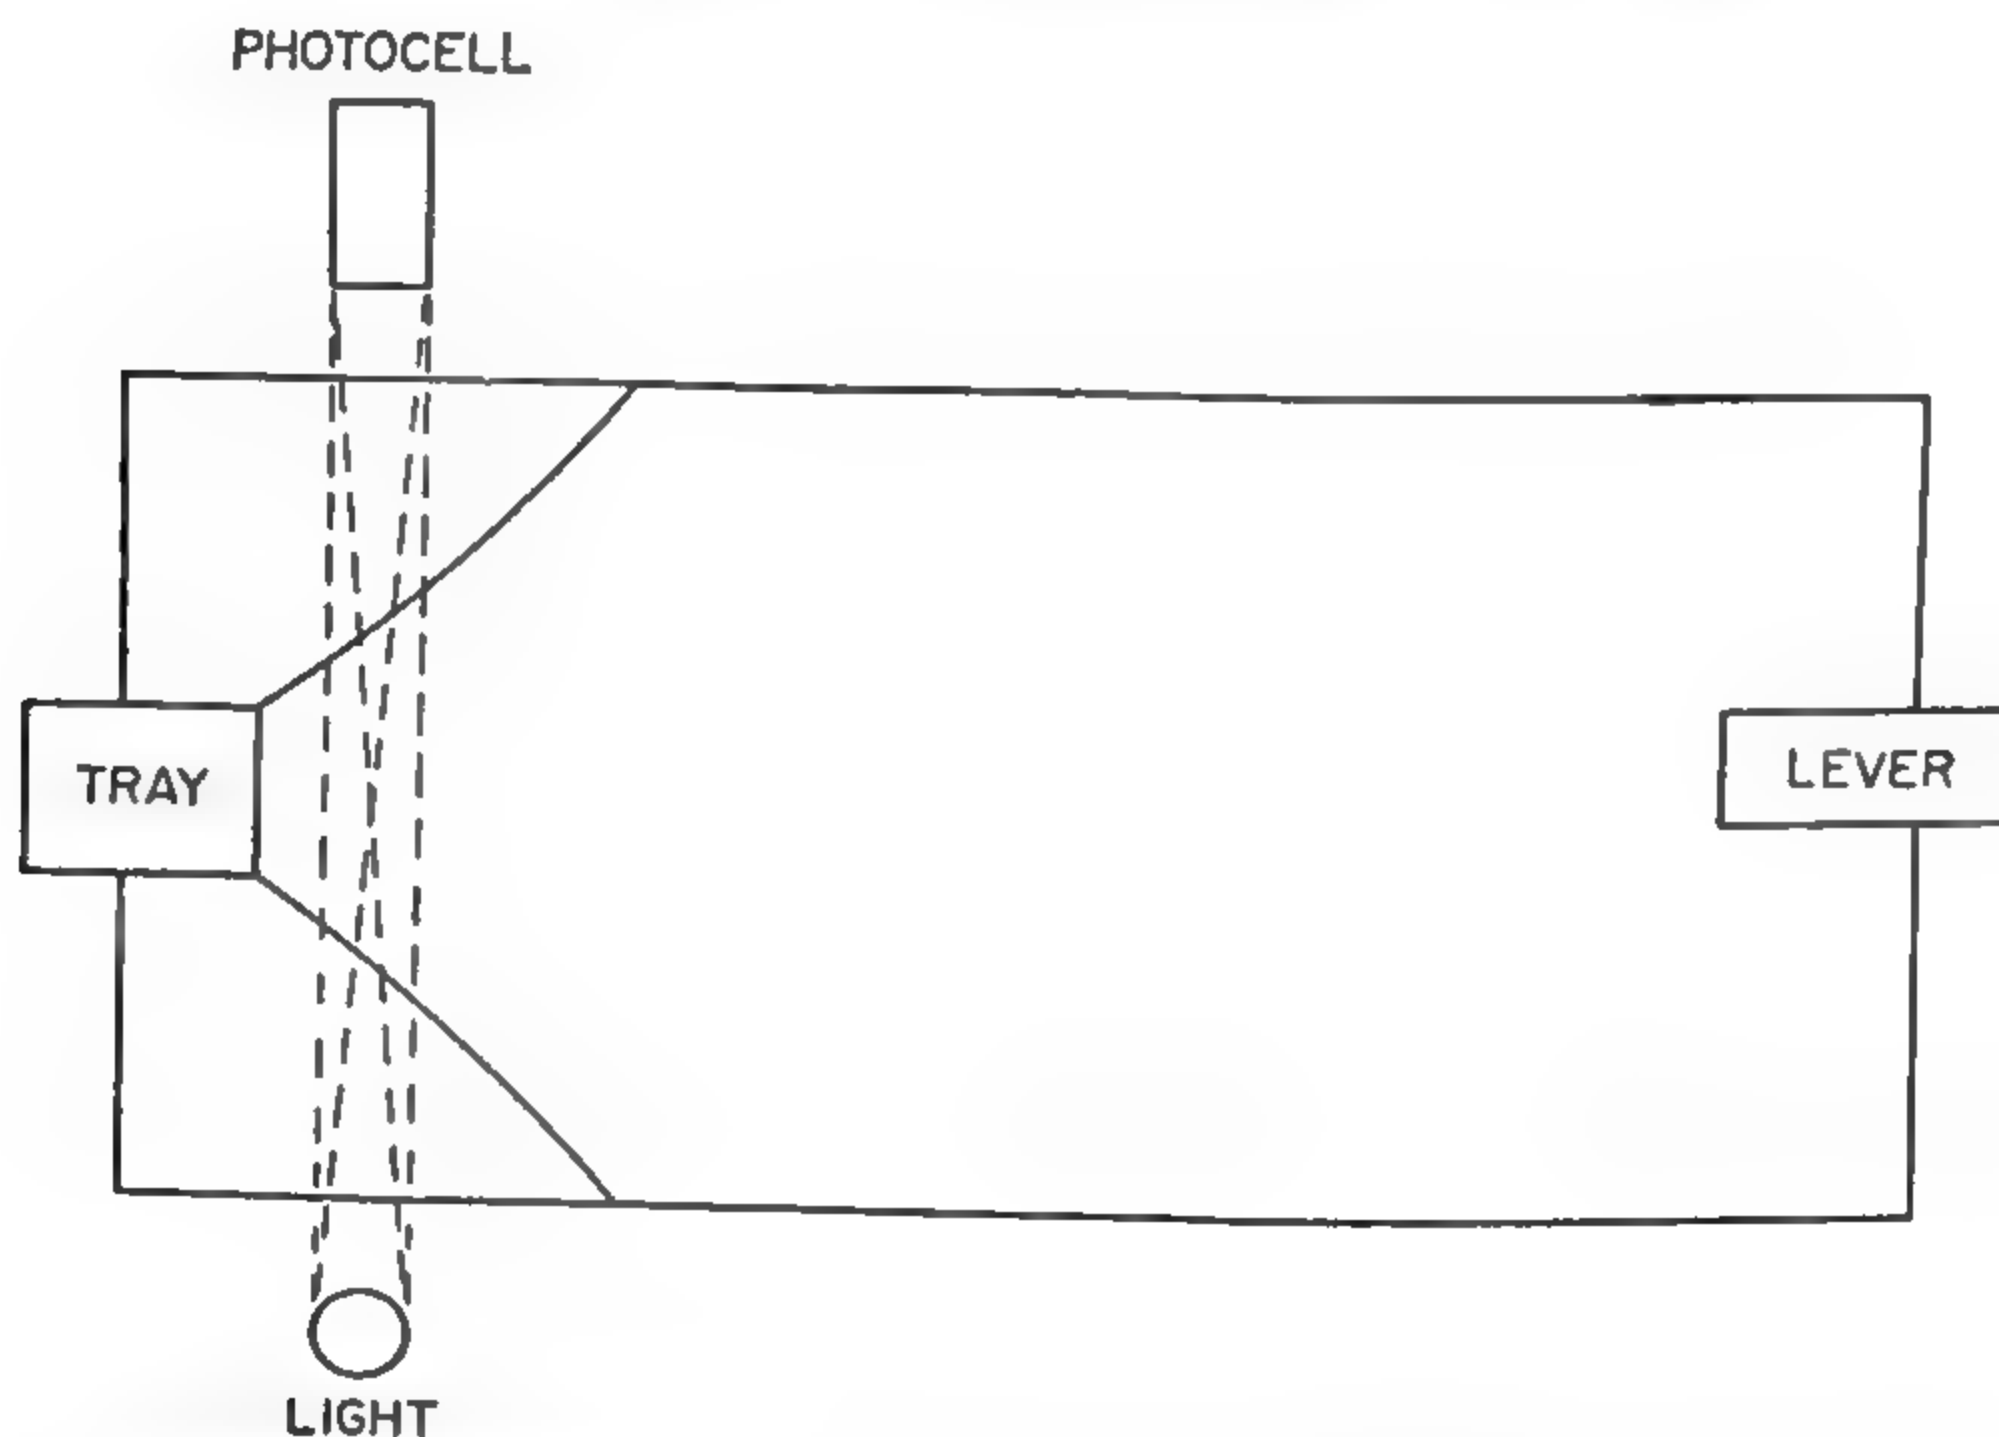


Figure 30. The floor plan of a Skinner box used to study sequences of lever pressing and eating. The lever and the food tray are on opposite sides of the box so that a clear-cut count of both responses can be made. Approaches to the food tray are counted by a photoelectric cell. (From Frick and Miller, 1951.)

tray is placed on the opposite side of the box from the lever; this makes the rat go back and forth across the box. The investigators tabulated the number of times the rat pressed the bar and the number of times it approached the food tray, as well as the order in which these responses occurred. Thus if A means approach to the lever and B means approach to the food tray, a typical sequence of responses might be as follows: AAAABBBBBBAAAABBAAAAAAABABBBAAA, etc.

These investigators studied the pattern of responses both before and after conditioning. Surprisingly enough, behavior was more stereotyped before conditioning than after. Stereotypy in this case usually meant repeating the same response in runs. In other words, if the animal approached the food tray (empty) before conditioning, it was more likely to repeat this act than it was to turn around and approach the lever. After conditioning, however, this pattern was broken. There was a tendency to alternate the responses (as perhaps the animals should) but this was less regular than the repetitive pattern before conditioning.

In another study, Frick (1953) applied this method of analyzing behavioral sequences to the problem of the gross effects of anxiety or fear on behavior. In this case, the patterns of A's and B's were studied by the *autocorrelation function*. The autocorrelation function tells to what extent later responses are dependent upon earlier responses and how many responses back this influence may extend. The results of Frick's study showed that a conditioned anxiety response induced by electric shock reduced the patterning in the lever-pressing-food-tray sequence without affecting the over-all rate of responding.

These studies represent initial steps in the study of the temporal organization of behavior. It is probable that we shall see many new applications of the basic technique of studying dependent probabilities in the near future. As we shall now see, these same techniques have been applied to problems in human behavior.

HUMAN VERBAL LEARNING

The most important and obvious problems in the serial learning of human beings occur in verbal behavior. Our verbal behavior consists of chains of words and phrases emitted in particular orders. These are clearly important, since the words do not readily make sense in a scrambled order. Try the following sequence on someone: "C will personnel by 20 applications room of new for building be positions in the officer received." The correct se-

quence is "Applications for new positions will be received by the personnel officer in room 20 of building C." Consequently, the learning and emission of verbal behavior in relatively fixed and invariant orders is important to our use of language.

There are other matters of great importance about verbal behavior, of course. Words carry meaning, and there are a variety of techniques for evaluating the role of meaning in verbal behavior, some of which we shall deal with later. Furthermore, it is clear that verbal behavior can be controlled and shaped by reinforcement in much the same way as the behavior of the rat in the Skinner box. Greenspoon (1955), for example, was able to control the rate of emission of plural nouns in an interviewlike situation by making plural nouns contingent upon some simple indicator of satisfaction on the part of the interviewer (saying "um-hum"). Other studies show that verbal behavior can be controlled by stimulus conditions and reinforcement schedules much as can any other instrumental activity (Buss, Weiner, and Buss, 1954; Verplanck, 1956). While this is important, the mere fact that verbal behavior is like any other class of instrumental behavior shows that this kind of control is not important to the unique properties of verbal responses—their chaining and their meaning. Let us therefore turn directly to the serial characteristic of verbal learning.

Methods in Rote Verbal Learning

In rote verbal learning we consider the influence of chaining of responses upon learning, performance, and retention. Ordinarily, when we are faced with the task of learning some new verbal responses such as memorizing a poem, we sit down with the material in front of us and, by a hit-and-miss procedure, learn it. We seldom memorize a poem by starting at the beginning and going through it line by line. More often we repeat lines in small groups or skip from one part to the next.

This hit-and-miss procedure will not do if one is interested in studying response organization, or chaining, in verbal learning. In laboratory studies of human learning, we are interested in the relationships between the parts of the material to be learned and how these relationships determine features of learning. In order to study these, it is necessary to hold constant the order in which the material is presented to the subject. Thus, in laboratory learning, the method of learning is partially controlled by presenting to the subject only one unit of material at a time in some predetermined order.

This technique reduces variability in methods of learning among subjects; that is, it makes our subjects learn more nearly in the same way.

Since the technique permits a control of serial relationships in the material to be learned, it also makes possible a systematic study of the influence these relationships have on learning. The serial-interaction effects are fundamental to verbal behavior, and it is essential that laboratory studies of human learning be well controlled in this respect.

Types of verbal material. As we shall see later, one of the most important determiners of rate of verbal learning is the *type* of verbal material learned, and it is important that this be carefully specified in every experiment. Furthermore, the verbal items in a given chain differ enormously from one another in important characteristics. Look, for example, at any of the sentences on this page. Such words as "in," "a," "the," etc., are obviously different from such words as "control," "material," "methods," etc. Consequently, seldom have investigators of verbal learning used ordinary sentences in natural languages as experimental material. Usually psychologists have made up their own verbal material so that some of the uncontrolled factors in ordinary sentences are eliminated or held constant.

One of the first investigators of human verbal learning, Ebbinghaus (1885), devised the most widely used technique for controlling the kind of verbal material used in such experiments. He showed that we can use the previously acquired verbal skills of human beings, and at the same time eliminate previous learning for the specific verbal material we wish to have the subjects learn. This can be accomplished with nonsense syllables. These are syllables novel enough to ensure that a uniform minimum of the subject's verbal skills is called upon. Indeed, if our syllables are truly "nonsense," the subject cannot bring into the experiment a large assortment of acquired reactions to the specific items he is to learn.

The nonsense syllable is usually constructed of three letters, most frequently two consonants with a vowel between them. Thus SYL or NEP is a nonsense syllable. Some of these three-letter nonsense syllables resemble English words more than others (for example, compare XYQ and CET). For this reason, several investigators have "calibrated" nonsense syllables in terms of their association value (Glaze, 1928; Hull, 1933). Such scaling has resulted in classifications of nonsense syllables according to the degree of meaning or familiarity. Incidentally, the association value, in terms of the number of associates given to nonsense syllables and words, has recently been itself used as a variable in verbal learning (Noble, 1952b; Mandler, 1955).

Other kinds of material, such as lists of adjectives, etc., have been calibrated in many ways. Melton (cited in Hilgard, 1951), for example, has

scaled lists of adjectives according to a measure of their degree of resemblance. This kind of calibration is important in the study of generalization or similarity in verbal learning. Various statistical tables (Thorndike and Lorge, 1944) giving the relative frequency of occurrence of words in ordinary English prose are an important source of controlled verbal material for the study of learning. The verbal material an investigator wishes to use in a particular study will depend upon the purposes of his study. There is now available a wide variety of material, and as new needs arise, new controlled material will be developed.

The serial method. In rote verbal learning the method of presenting

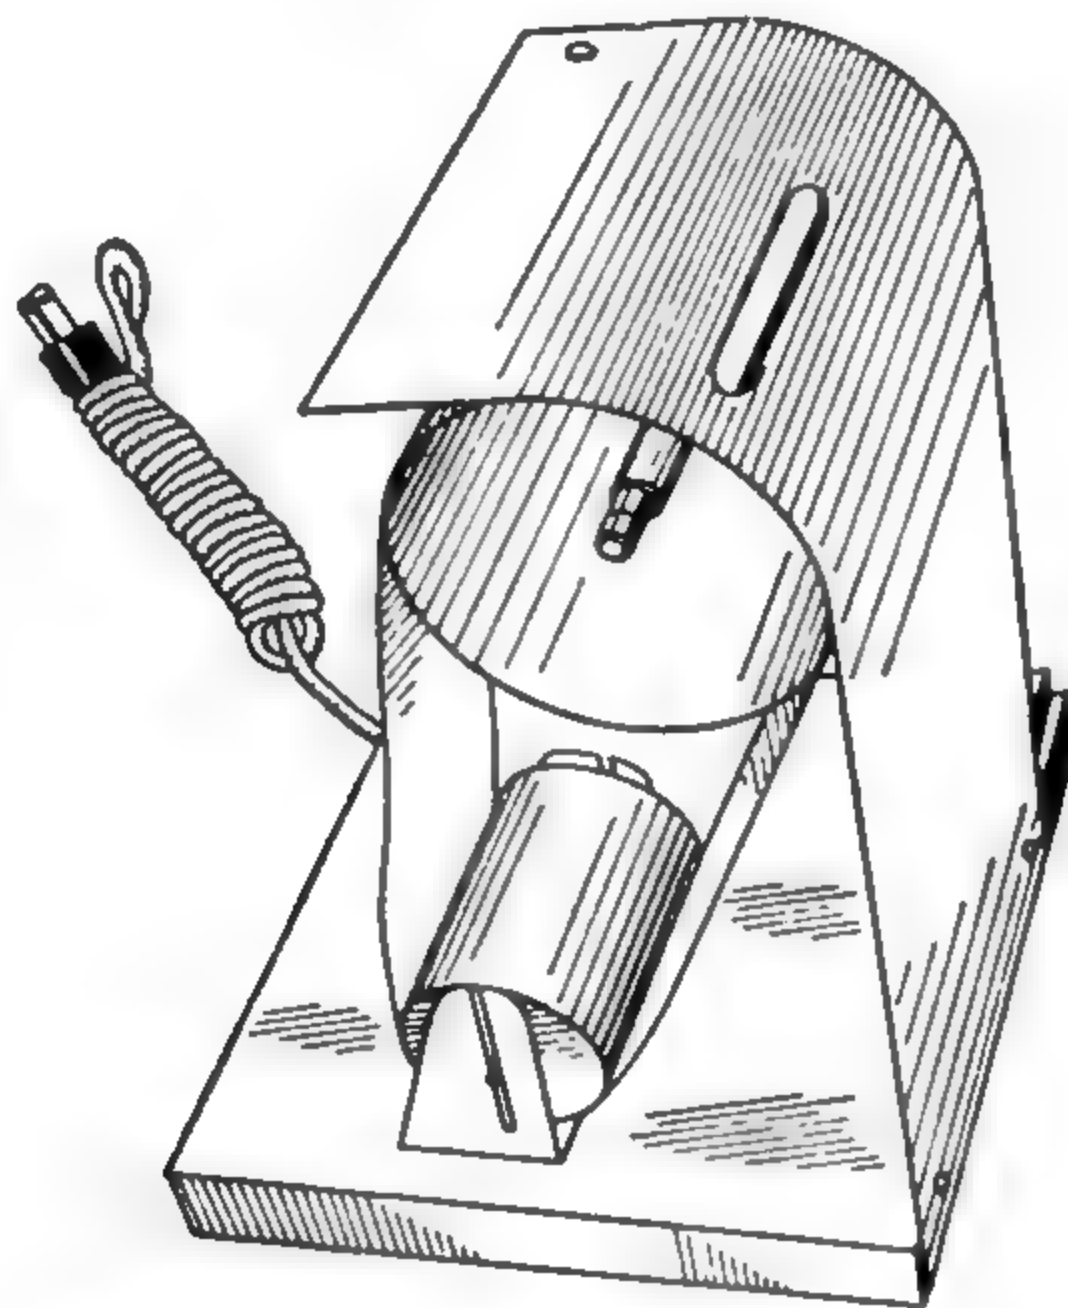


Figure 31. A memory drum. The syllables or words are printed on the paper behind the shield. The drum does not move continuously but moves each item into place and then waits a brief time before moving the next one into place. (Courtesy of Ralph Gerbrands.)

the material is usually controlled in order to force certain kinds of relationships between the items in the material. In the *serial method* the material is forced into a rigid sequence. A memory drum (see Figure 31) is usually employed to achieve this order. The items to be learned (nonsense syllables, for example) are exposed to the subject one at a time in the small window in the memory drum. These items are usually exposed at a constant rate, so that the same amount of time is spent on each item. The subject is instructed to *anticipate* each syllable before it appears in the window. There is sometimes a blank space between each syllable, which allows the subject time to anticipate. Usually the subject is instructed to spell rather than say the syllables, if nonsense syllables are used, in order to avoid pronunciation difficulties.

In the serial method, each item serves in turn as a stimulus and as a response. For example, the first item in a list might be the nonsense syllable SYJ. When the list begins on each trial, the stimulus for this syllable is provided by the symbol used to indicate the beginning of the list. The next syllable might be LOZ. SYJ is the stimulus which tells the subject that LOZ is coming up. LOZ in turn might be followed by NEP. Thus, SYJ is the stimulus for the response "L-O-Z." LOZ is the stimulus in turn for the response "N-E-P."

Table 1 shows the results of an experiment on rote verbal learning with one subject. In the columns are given the 12 syllables in the list. The rows represent the trials, and there were 20 given to this subject. Note that a trial is defined as one repetition of the entire list. The plus signs represent correct anticipations, and the zeros represent incorrect responses or no responses at all. In some experiments, incidentally, we may be interested in recording the exact errors made by the subject. The first time through the list, of course, the subject can get no responses right, since this is the first time he has seen the material. Notice that this particular subject guessed at the first syllable on the second trial and got it right. On the third trial he correctly anticipated the first syllable, the second syllable, and the last syllable. On the nineteenth and twentieth trials the subject anticipated all the syllables correctly.

If we are interested in studying the retention of verbal material, some time after the original learning we may give a test of retention. Such a test could be obtained in exactly the same way the learning data were obtained; the subject is required to anticipate each syllable in turn. On other occasions, we may be interested in a free-recall test; the subject is simply asked to recall as many of the syllables as possible in their correct order.

The paired-associates method. For many purposes there is one grave defect in the serial method—each item serves as a stimulus and a response in turn. Sometimes it is important to sort out and independently vary the stimulus and response aspects of a task to be learned. This cannot be done with the serial method; consequently, a modification known as the *method of paired associates* has been developed.

In this method, as its name implies, pairs of items are presented to the learner. He is instructed to learn the pairs in such a way that when the first member of the pair appears, he is able to recall the second. Thus the first member of a pair is a stimulus item and the second member a response item. Unfortunately, in this method, the serial order of presentation cannot be retained, since if it were, subjects could make a connection between re-

TABLE 1
Subject's Record Sheet from Serial-learning Experiment

Trial	Syllable											Number	
	CEX	MOQ	RUY	GAF	LIQ	KOC	QUZ	DEJ	TAH	WOG	FIK	VUS	correct
1	0	0	0	0	0	0	0	0	0	0	0	0	0
2	+	0	0	0	0	0	0	0	0	0	+	+	3
3	+	+	0	0	0	0	0	0	0	0	0	+	3
4	+	+	0	+	0	0	0	0	0	0	+	0	4
5	+	+	0	+	0	0	0	0	0	0	+	+	4
6	+	+	+	+	0	0	0	0	0	0	+	0	5
7	+	+	+	0	0	0	0	0	0	0	+	+	5
8	+	+	+	0	0	0	0	0	0	0	+	+	5
9	+	+	+	+	0	0	0	0	0	0	+	+	6
10	+	+	+	+	0	0	0	+	0	+	+	+	8
11	+	+	+	+	0	0	0	0	0	+	0	+	5
12	+	+	+	+	0	0	0	0	0	+	+	+	8
13	+	+	+	+	0	0	0	+	0	+	+	+	8
14	+	+	+	+	+	0	0	0	0	+	+	+	8
15	+	+	+	0	+	0	0	0	+	0	+	+	7
16	+	+	+	+	+	+	+	0	+	0	+	+	10
17	+	+	+	+	0	+	+	0	+	+	+	+	10
18	+	+	+	+	+	+	0	+	+	+	+	+	11
19	+	+	+	+	+	+	+	+	+	+	+	+	12
20	+	+	+	+	+	+	+	+	+	+	+	+	12
Number correct	19	18	15	14	6	5	4	5	6	8	16	17	

sponse items and the stimulus items which follow. Therefore, in this method, the pairs of stimulus and response items are presented in some random or prearranged order which is different for every trial. Obviously, the method of paired associates destroys the simple serial characteristics of rote learning, but it enables the investigator to distinguish between stimulus material and response material. For some problems the simple anticipation method is more suitable; for others the method of paired associates provides the most useful technique. The uses of this method will be described in later chapters.

Associations in Rote Verbal Learning

One of the fundamental problems is the degree to which different items in a verbal-learning task are associated with one another. In the method of serial learning we know that successive items are associated, since otherwise the learner would not be able to anticipate. It turns out, however, that learners also make connections between items that do not adjoin. These are called *remote associations*.

The nature of remote associations. In a typical experiment, nonsense syllable A becomes a cue or stimulus for the response of nonsense syllable B. B is a cue for the response C, etc. The subject learns the list in order. In addition, however, syllable A is connected to some degree with syllable C, so that A will occasionally give rise to the response C. Such responses, of course, are analogous to anticipatory errors found in the studies of the linear maze. In verbal learning these anticipatory responses are called *remote forward associations*.

Since perseverative errors are found in the linear maze, it might be supposed that the same errors would be found in rote verbal learning as well. Indeed, they are. If the learner responds with A when he should have responded with D, a perseverative error has occurred. In verbal learning, such errors are called *remote backward associations*.

Ebbinghaus (1885) was the first to detect these remote associations through a technique known as the *method of derived lists*. First he learned a list of nonsense syllables—Ebbinghaus always used himself alone as a subject—then learned a second list obtained by systematically rearranging the first list. Sometimes in the rearrangement every other item was made adjacent, sometimes every third item, and sometimes every fourth. Ebbinghaus's results showed that he could relearn systematically rearranged lists much more quickly than he could relearn ones haphazardly rearranged or an entirely new list. Furthermore the fewer the number of items skipped in the rearranged lists, the quicker the time of learning. From this Ebbinghaus

inferred that remote (forward) associations had helped him in the learning of the rearranged lists, and therefore that these associations must have been learned in the original lists.

Ebbinghaus's results tell us little about the specific remote associations formed. Nor finally, does his method indicate the frequency of remote associations.

A method used by McGeoch (1936) tells us a good deal more about the exact nature of remote associations. He asked subjects to practice, by the method of anticipation, memorizing a series of nonsense syllables. He stopped the practice before the subjects had learned to anticipate perfectly

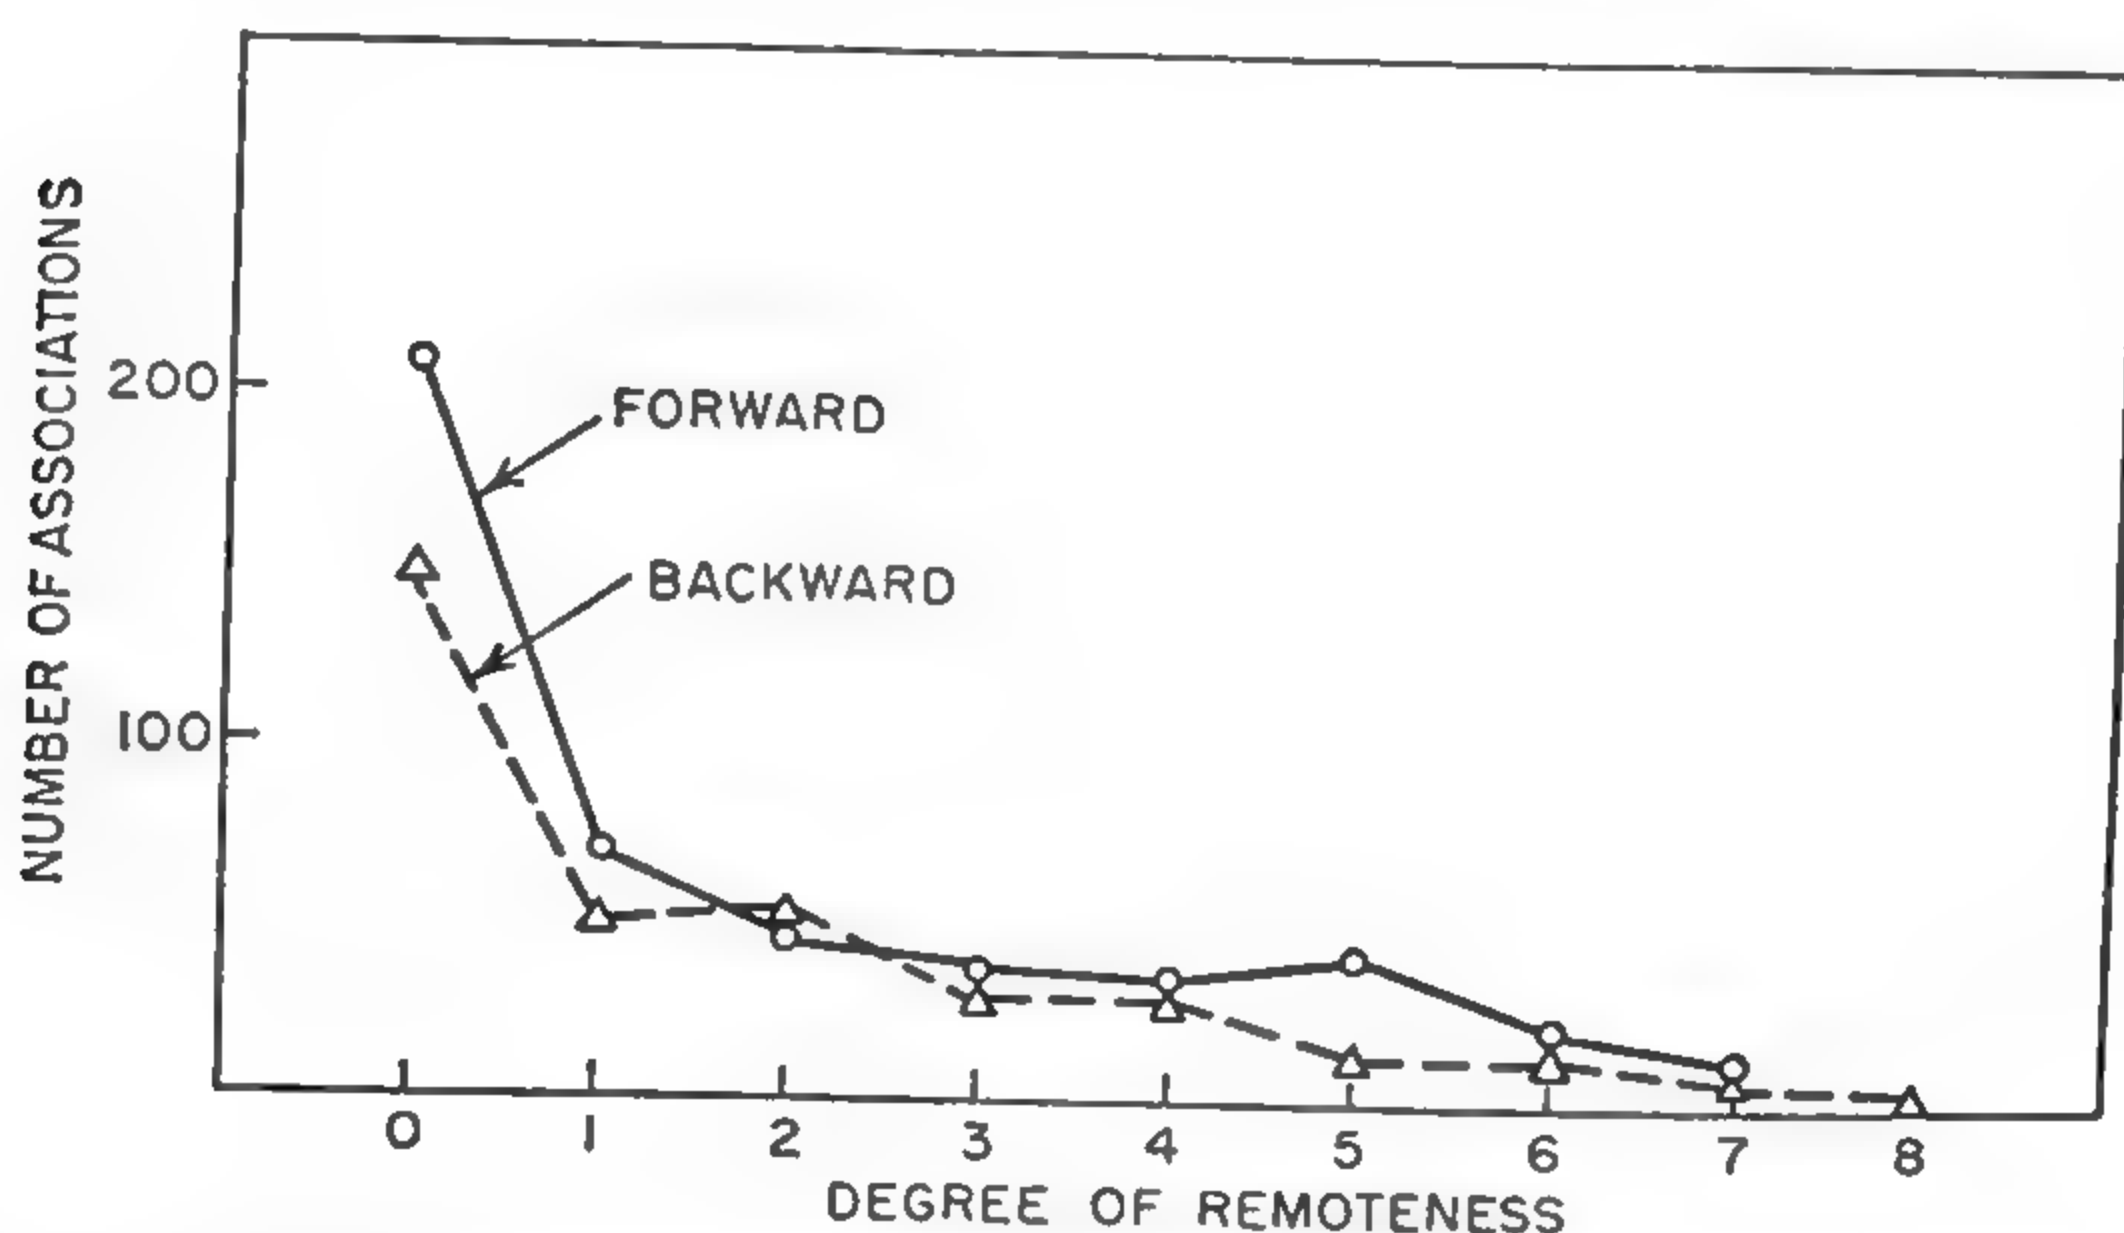


Figure 32. Number of associations as a function of the degree of remoteness (syllables spanned). (Data from McGeoch, 1936.)

and then gave them a free-association test with the nonsense syllables they had been practicing as stimuli. That is, he read in a scrambled order the nonsense syllables the subjects had been learning and asked them to respond with the first syllable which came to mind. Figure 32 shows the frequency of association as a function of degree of remoteness. Notice that the number of remote associations, both forward and backward, falls off abruptly after zero degrees of remoteness; in addition, it appears that forward associations are slightly more numerous than backward associations, at least at the first degree of remoteness.

Interpretations of remote associations. We shall deal with the significance of remote associations later, but now let us look at some of the notions about how remote associations are brought about. One idea is that the earlier items are still being responded to when the later items appear. Thus

there is an association between a trace of the earlier item and the occurrence of the later one. This theory has been greatly elaborated by Hull and his associates (Hull, et al., 1940) and used to explain many characteristic phenomena of rote serial learning.

Another theory interprets remote associations as mediated connections between two items in a list and a third item common to them. The mediating factor might be something within the subject himself or some constant feature of the environment (the context) to which the specific items on the list become connected (McGeoch, 1942). There is experimental evidence that mediated associations can produce learning. Suppose, in a paired-associate experiment, we ask subjects to associate nonsense syllables to common words, so that given the nonsense syllable, the subject can recall the appropriate word. For example, given CEF, the subject would learn to say "stem." Now "stem" is a strong eliciting stimulus, in a free-association test, for the word "flower." And "flower," in turn, is a strong eliciting stimulus for "smell," but "stem" does not tend to elicit "smell" directly. If we now take a series of such chains and ask the subjects to learn the third member of the chain (smell) to the original nonsense syllables, apparently they learn more rapidly than they would learn a list of unrelated words (Russell and Storms, 1955). In other words, the associative habits built into the subjects before the experiment result in the mediation, or transfer, of the pairing of the same nonsense syllables from one list of words to another. This, of course, tells us nothing about how such mediational activity comes about, but it does suggest that associations between two items arise from the mediation of chains of intervening associations.

Remote forward and backward associations in rote verbal learning may have more than a superficial similarity to anticipatory and perseverative errors in linear maze learning. Anticipatory and perseverative errors are attributable to stimulus generalization. It is difficult, however, to see how simple gradients of stimulus generalization in verbal learning could be based exclusively upon resemblances among the items themselves. Such similarities would not always produce gradients of remote associations, since the resemblances might be distributed in a chance manner throughout a list. It is possible, however, that the context in which a particular item is to be emitted allows for stimulus generalization (see Schulz, 1955). If, for example, in a list of nonsense syllables DOX is followed by CET, DOX and the context provided by having gotten through so many items in the list elicit CET. Suppose, however, that the context generalizes to some later point in the list. DOX might then elicit some other response, say ZYM. In intuitive terms this would happen

when the subject had "lost his place" in the list. It is even possible that such contextual generalization might be aided by random resemblances among the syllables themselves, as perhaps would be the case if DOQ preceded ZYM in the list.

If the notion that contextual and item generalizations are responsible for remote forward and backward associations is correct, then the number of remote associations ought to increase directly as the similarity between items in a list is increased. Although there is no direct evidence that this is so, there is some indirect evidence. Underwood and Richardson (1956) find, for example, that rate of learning is a function of the similarity of items within a list. The higher the similarity, the lower the rate of learning. This means that the greater the generalization between items on a list the slower the learning, and, by implication, this lower rate of learning is the result of a greater number of remote associations intruding at the wrong point in the list.

The Serial-position Effect

For a long time students of learning have been aware that there must be an intimate connection between remote associations and other characteristic phenomena of rote verbal learning. One of these is the *serial-position effect*. This effect refers to the fact that the rate at which a given item in a serial list of items is learned depends upon the location of that item within the list. When learning is by the serial method, the beginning of a list of nonsense syllables is learned most quickly and the end of the list is next easiest to learn. For this reason, the serial-position effect has sometimes been called the *primacy-finality* effect, to emphasize this point.

Serial position in rote serial learning. We can see what the serial-position effect looks like by examining the curve in Figure 33. This curve is taken from an experiment on the memorization of nonsense syllables by the serial method (Hovland, 1938b). The total number of errors or failures of anticipation at each position in the list is shown on this curve. The greatest number of errors occurred at the seventh syllable position, which is just past the middle of the list.

The shape of this curve has wide generality. The most difficult part of a list of serial items learned by anticipation is just past the middle of the list. As more items are added, the point of maximum errors always remains in the same relative position (Hovland, 1940). The serial-position curve is essentially the same for meaningful words as it is for nonsense syllables, though, of course, the total number of errors is less in the former case (P. c.

Crary and Hunter, 1953). If one compares early performance with late performance in learning, of course, the serial-position curve is flatter in the later case, but this is simply because all of the items are finally learned (Ward, 1937).

Explanation of the serial-position effect. The best-known attempt to explain the serial-position effect for anticipatory rote learning is that of Hull (1935). Hull argued that remote forward associations *inhibited* immediate associations between them. For example, a remote association be-

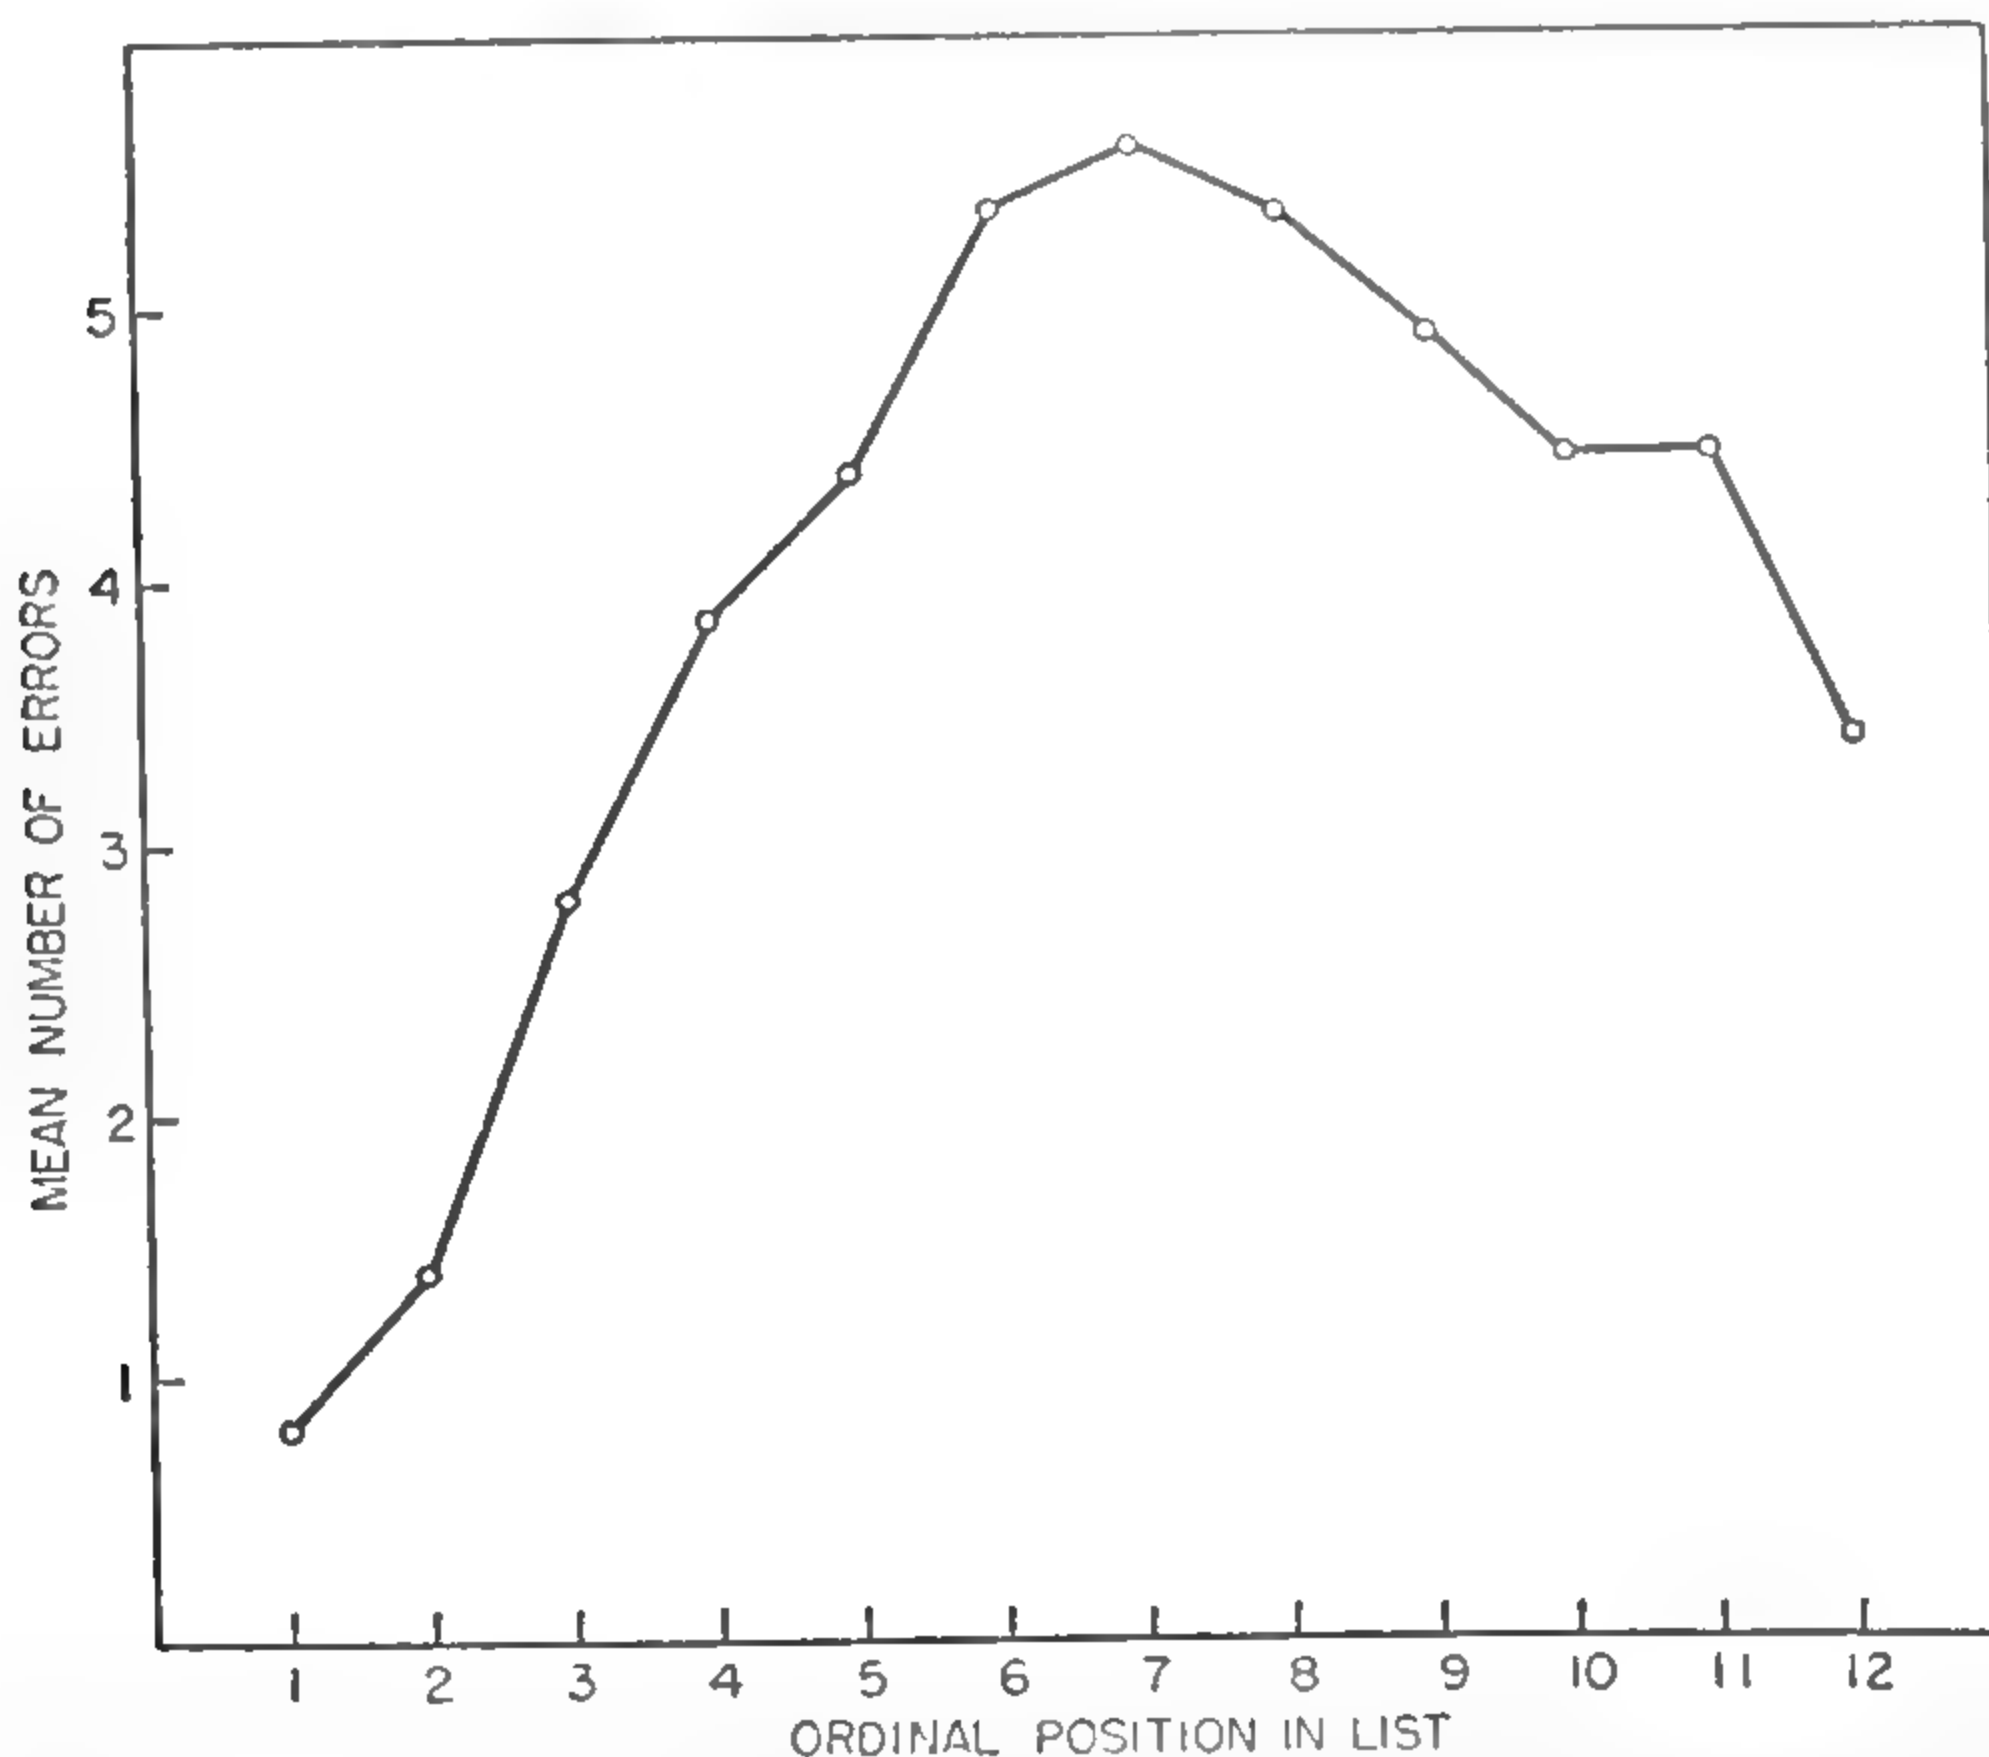


Figure 33. The number of errors at each position in a list of nonsense syllables during the learning of the list by serial anticipation. (Hovland, 1938.)

tween items 1 and 9 would inhibit or tend to reduce the tendency for all intervening direct associations to occur. Consequently, the items in the middle of the list should be inhibited in proportion to the number of delayed associations which span them. This notion is illustrated in Figure 34 (Hull, 1935). Notice that the greatest number of delayed associations bridges the middle items, and this is a way of representing the notion that the most difficult items to learn should be in the middle.

There are two things wrong with Hull's hypothesis. First of all, it ignores remote backward associations; this is important because remote forward and backward associations are not perfectly symmetrical. Secondly, it ig-

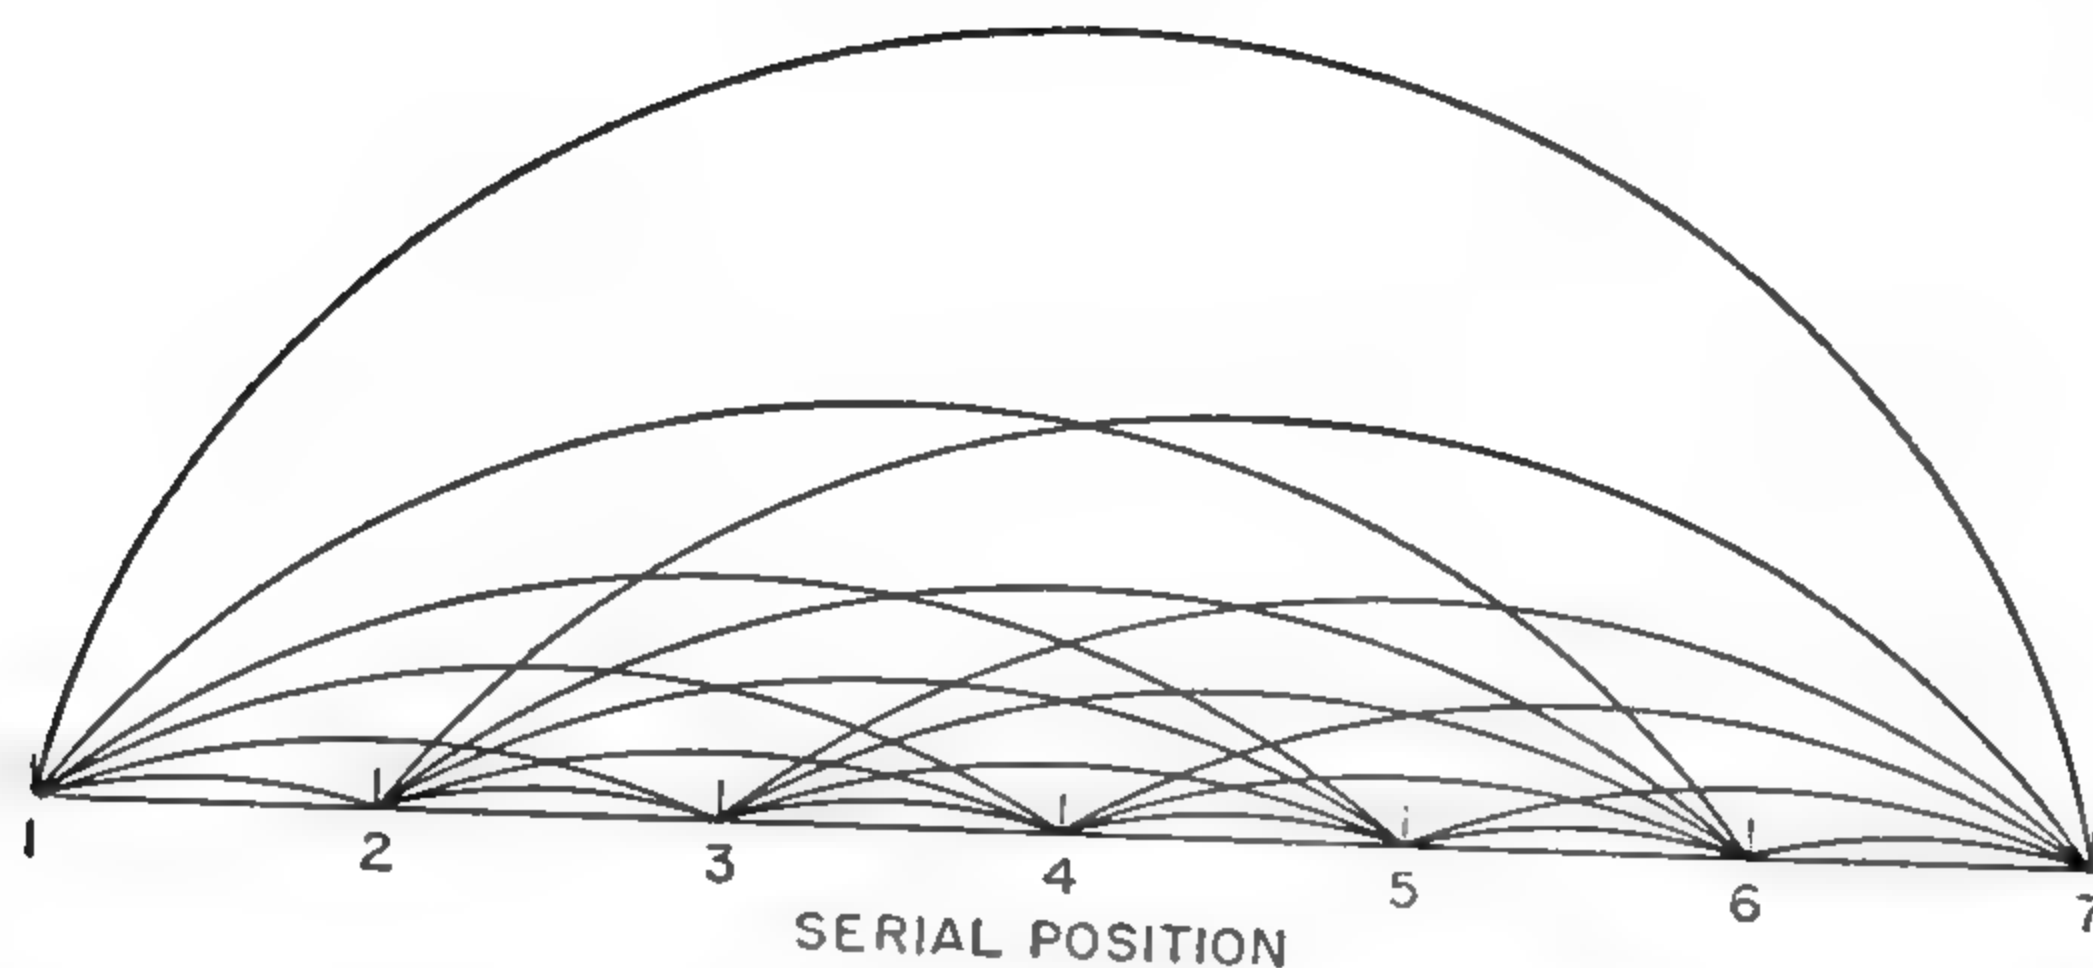


Figure 34. Illustrates how remote forward associations could account for a symmetrical serial-position curve. The number of errors or omissions at each position is proportional to the number of associations spanning that portion. (Hull, 1935.)

nores the fact that the middle item is not the most difficult to learn in actual practice; the items just past the middle are. Hull and his associates (Hull, et al., 1940) tried to remedy this difficulty in a general theory of rote learning, but the results were so terribly complicated that they have been regarded as rather unconvincing.

Bugelski (1950) pointed out that remote associations could much more convincingly account for the serial-position effect if we simply assume that they interfere with the emission of correct responses. This would place the greatest interference at slightly past the center of the list. Thus, while Bugelski's analysis produces some skewing in the theoretical serial-position curve, it does not produce as much as actually exists in practice.

If interference from remote associations is the primary factor in producing the serial-position effect, it ought to be possible to test this directly by counting the intrusions of wrong responses at each position. Unfortunately, the usual method of presenting material in serial-anticipation experiments does not permit this to be done. The syllables are paced at a moderately fast rate, so most of the "errors" are simply failures of response, rather than intrusions. This may be because the subject takes too long in getting the erroneous response out, and by the time he is ready to do so, the next syllable has appeared. Deese and Kresse (1952) performed an experiment in which they allowed unlimited time for the subject to respond between syllables. When this is done, it appears that the remote associations (intrusions of other items from the list) are almost symmetrically distributed about the middle of the list. The displacement of maximum difficulty from the middle to just past the middle is explained by the fact that, as the list proceeds, the subjects are simply unable to respond with even a reasonable guess. Thus it

appears as though the serial-position effect is composed of two components, an almost symmetrical component of overt intrusions or remote associations, and an increasing tendency through the list to be unable even to guess.

Though the tendency to make intrusive errors based upon remote associations accounts for the bowed shape of the serial-position curve for rote serial learning, it does not account for the skewing, or displacement, of the curve. This displacement is probably the result of a combination of several factors including memory span and perhaps something like the reactive inhibition factor discussed in the chapter on extinction.

The serial-position curve in free recall. Suppose that a list of nonsense syllables or randomly arranged words is presented to a subject, and, instead of recall by serial anticipation, the subject is allowed to recall the words in any order he chooses. Under these conditions the most frequently recalled words are at the end of the list (Welch and Burnett, 1924; Deese and Kaufman, 1957). Two sample serial-position curves for such free recall are presented in Figure 35. (Note that these curves are plotted in correct responses rather than errors.) Thus it appears that the serial-position curves for serial anticipation and for free recall are roughly the reverse of one another. In

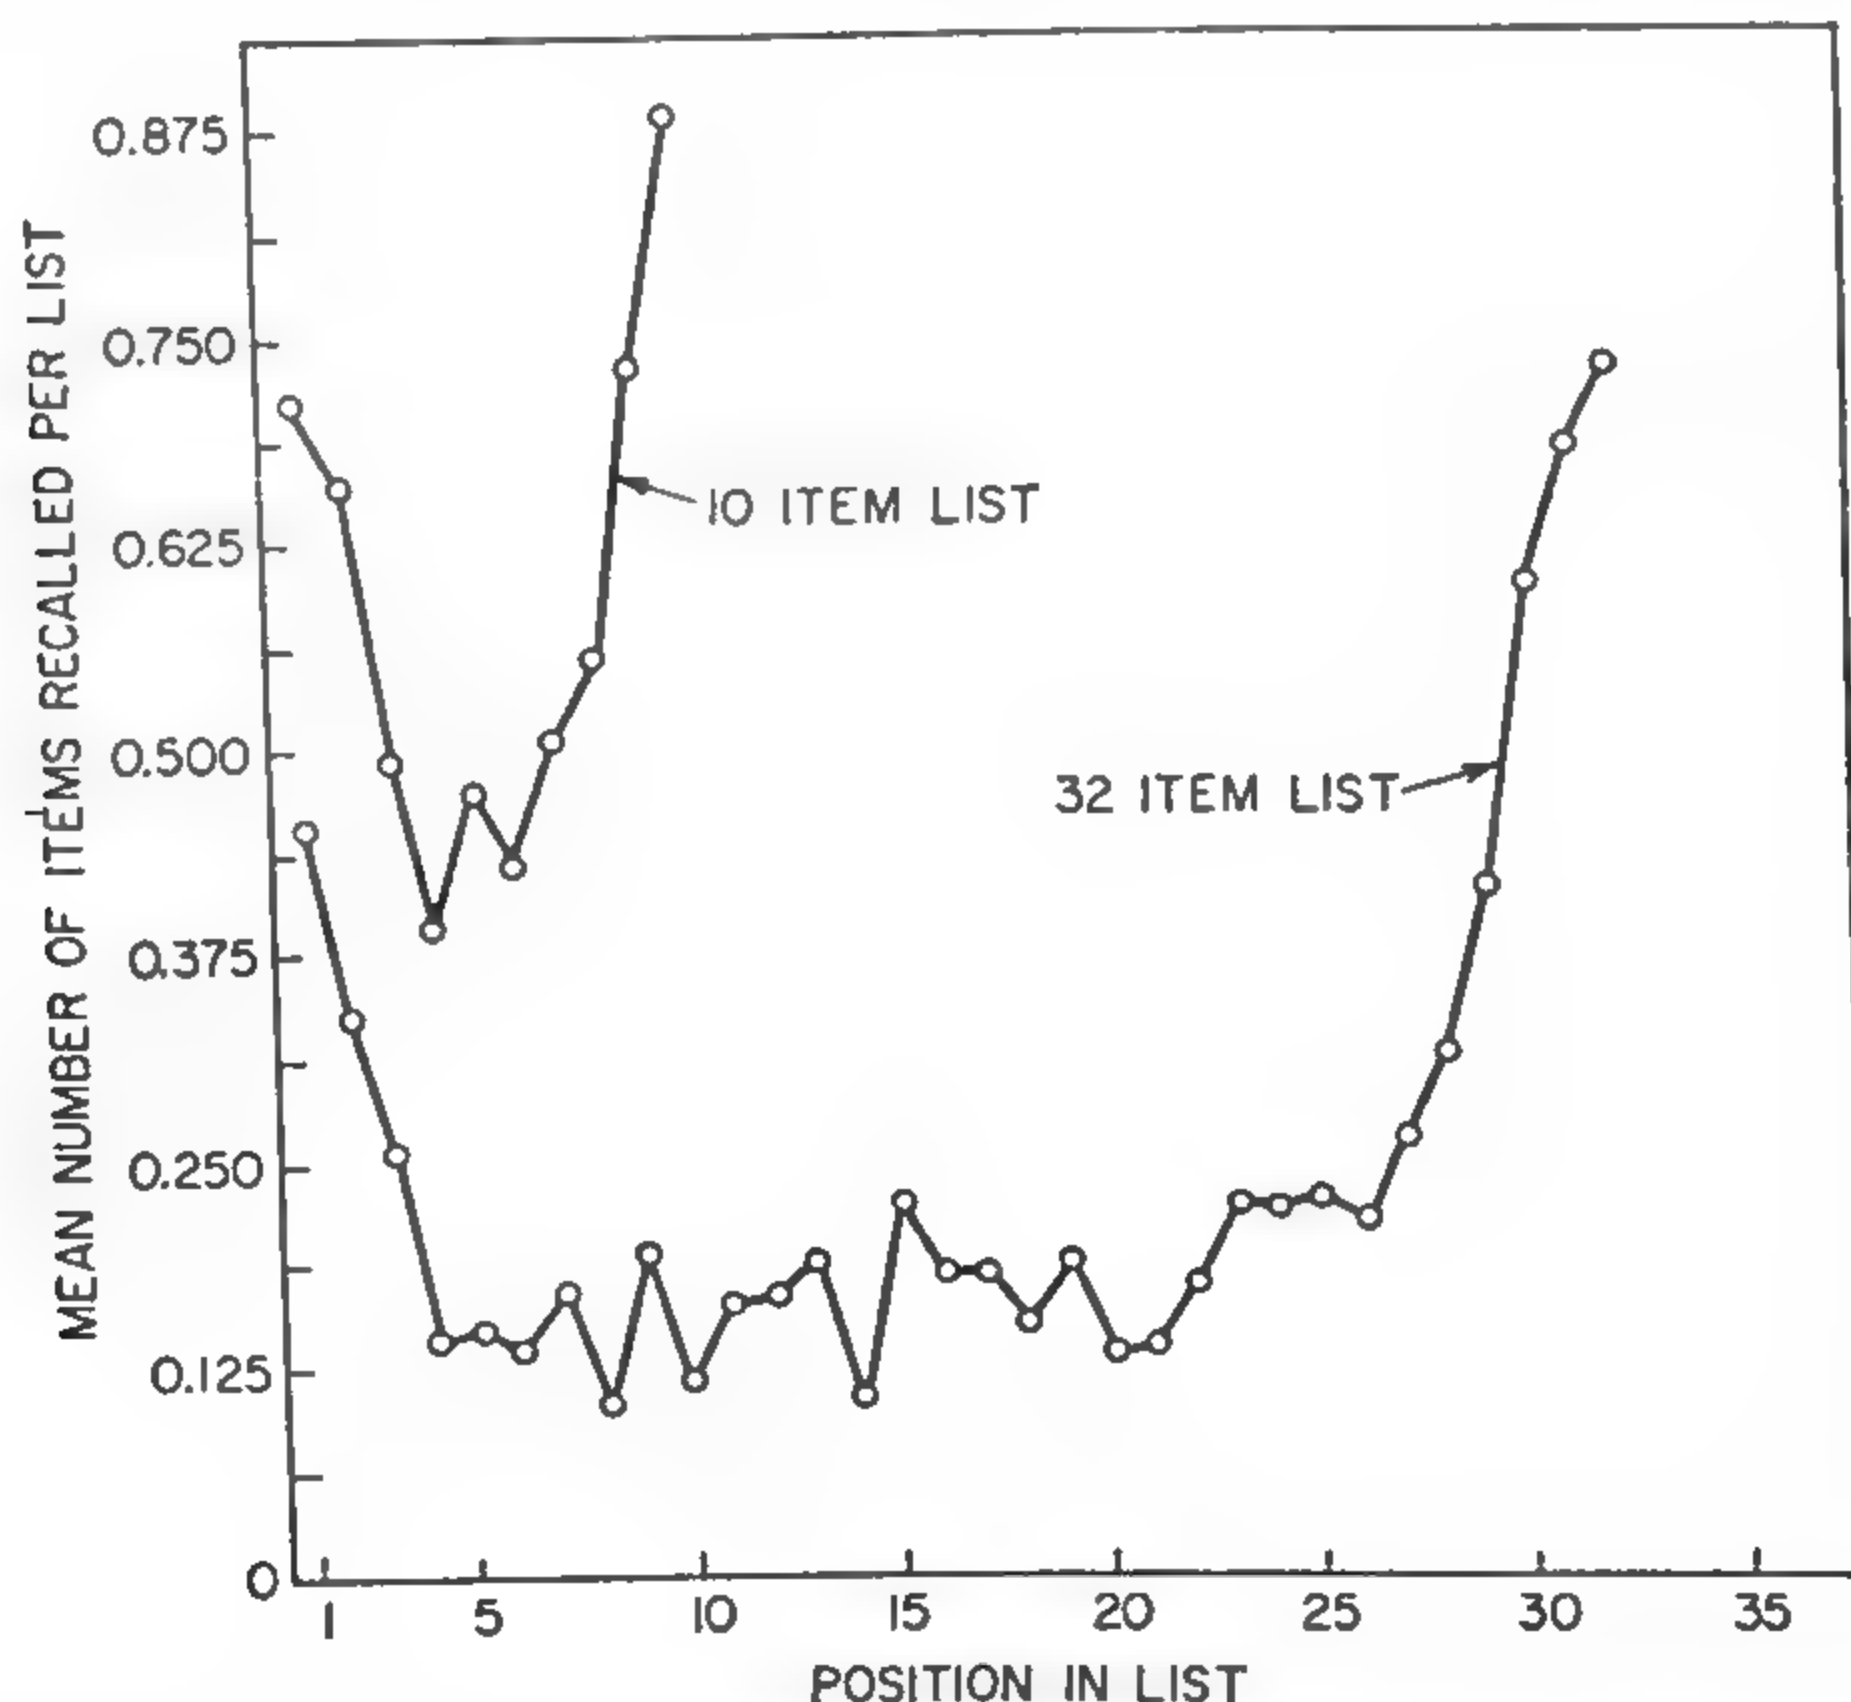


Figure 35. Serial-position curves for free recall. In contrast to the curve in Figure 34, these curves are presented in terms of the number of items correctly recalled. The highest frequency of recall is at the end of the list. (Deese and Kaufman, 1957.)

one, the greatest number of correctly recalled items is at the beginning of the list, and in the other, the greatest number is at the end.

It turns out that in free recall the greatest determinant of the probability of recall (as well as the order of items in recall) is the associative strength of the item (Bousfield, Cohen, and Silva, 1956). Thus in the immediate free recall of randomly arranged items it is probable that the last few items in the list are strongest, perhaps because of the persistence of an image of them; therefore, these items are recalled first and have the greatest probability of being recalled.

Something quite different seems to happen with the free recall of connected prose, however. If brief paragraphs of English are used as learning material and recall is unrestricted, the resulting serial-position curve looks more like that of the method of anticipation rather than that of the free recall of random items (Deese and Kaufman, 1957). This is, of course, because the recall of connected material is not really free; the serial characteristics of ordinary prose force the subject to start from the beginning and go through to the end, at least if he is to take advantage of the redundancy (see below) of language to maximize his recall. Thus, free recall of unorganized material is markedly different from that of highly organized material. A systematic exploration of this difference shows that as more and more organization is introduced into a list of words, the characteristic of the serial-position curve changes gradually from that of the free recall of randomly arranged items to that of recall by serial anticipation (Deese and Kaufman, 1957). This result seriously suggests that the method of serial anticipation is much more like the recall processes used for ordinary linguistic material than we had heretofore supposed.

Statistical Organization of Verbal Material

In the last few paragraphs we have discussed the question of organization in verbal material without really saying what we mean by "organization." In this section we shall discuss the meaning of verbal organization of a particular kind—statistical organization—and how it affects the sequential learning of verbal material. Statistical organization involves taking account of dependent probabilities between items in exactly the same way that dependent probabilities have been applied to multiple-choice behavior. Since earlier in this chapter we explored this problem, we are already familiar with the general concepts.

A moment's thought will make it clear that dependent probabilities apply to the elements of languages. Take the occurrence of letters in written Eng-

SERIAL LEARNING

lish, for example. Everybody knows from Poe's *The Gold Bug* that letters in English do not occur with equal frequency. The letter *e* appears more frequently than any other letter; consequently, if we were to guess that the letter *e* occurs in a given sequence of letters rather than simply picking a letter at random, we should have a better chance of being right. Likewise, the occurrence of a given letter is influenced by what has gone before. Therefore we can use our knowledge of what has preceded in a particular passage to predict what is going to happen next. In English, the letter *q* is always followed by *u*; thus the letter immediately after *q* is perfectly predictable. It is, to use the term frequently used in such analyses, an example of *perfect redundancy*—the *u* tells us nothing new, and if it weren't for habitual usage, might just as well be left out. There are associations between all other letter combinations in English, though not quite so strong as that between *q* and *u*. The letter *t*, for example, is more likely to be followed by *h* than by *m*.

It is easy to see, furthermore, that the probability of occurrence of a given letter will be influenced by the letter which precedes it by two or three spaces and more. In English, for example, two consonants are more likely to be followed by a vowel. The probability of occurrence of a given letter in English, as a matter of fact, is influenced to a significant degree by letters as many as seven positions in front of it (Newman and Gerstman, 1952).

The method of actually calculating the probabilities involved in such organized sequences of events in language derives from the mathematical theory of communication or information. *Information theory*, as it is usually called, has had important methodological implications for the study of human language and learning. We shall not deal with the theory directly in this book, although it will indirectly enter into many of the problems of language learning that we shall face. A general introduction to information theory is given in Miller's *Language and Communication* (1951). For now, however, let us simply accept the fact of dependent probabilities and the resulting statistical organization in language, and examine some of its implications in the serial learning of linguistic material.

Statistical organization and immediate free recall. A textbook on the psychology of learning is hardly needed to tell one that nonsense material is harder to learn and recall than connected, logical material. Some experimental studies show us, however, that it is not necessarily the sensibleness of connected material that is primarily responsible for this.

Miller and Selfridge (1950) show how gradual approximation to the statistical structure of ordinary language influences the amount of material recalled. An order of approximation to sensible English may be defined by

the number of preceding items which determine the present item, and in practice it serves as a way of constructing verbal material with various levels of dependent probabilities without going to the enormous labor of computing the necessary probabilities involved. If we select words at random from the dictionary or a word list, we should have a *zero-order approximation* to English. An example of a zero-order sequence might be: "pack byway phosphor kettle hypermetrical gloat." A *first-order approximation* would select words according to their relative frequencies in English usage. Thus common words like "the" and "before" would most likely show up. If we gave an individual a word to begin a sentence with and gave the second word of his sentence to another individual for him to start a sentence with, etc., we should be able to construct a *second-order approximation* from the resulting list of words. An example of a second-order approximation taken from Miller and Selfridge's paper looks like this: "was he went to the newspaper is in deep and." For *higher-order approximations*, individuals may be given sequences of words to use. In these cases, they add a word directly after the sequence and drop the first word of the sequence. This new sequence is then given to another individual, etc. The words immediately following the sequence of given words provide us with a higher-order list. An example of a seventh-order approximation obtained in this manner looks like this: "then go ahead and do it if possible while I make an appointment I want to skip very much around the tree and back home again to eat dinner after the movie early so that we could get lunch because we liked her method for sewing blouses and shirts is."

We can see that the seventh-order approximation appears to make vague sense—a sort of stream-of-consciousness sense. This is because it preserves the ordinary dependencies of speech without having anything in particular to say. Miller and Selfridge showed that it is this chain that determines how much material a person remembers. Up to a point, the higher the order of dependency, the easier it is to remember. With adult subjects, Miller and Selfridge found little or no difference in the recall scores for a seventh-order approximation and good, sensible English. Thus, even though the seventh-order approximation is essentially nonsense, it is recalled as well as the sensible material in a test of immediate free recall.

A summary of Miller and Selfridge's results is given in Figure 36. This shows the relationship between order of approximation and per cent of words correctly recalled. Notice that above fifth-order approximations there seems to be little improvement in recall. This apparently is because this and higher approximations preserve the short-range dependencies which

the individual can remember. Thus it appears that to some extent meaningful material is more easily learned because it preserves the short-range associations of ordinary language.

As we mentioned earlier, introducing statistical organization also changes the form of the serial-position curve. Deese and Kaufman (1957) studied the serial-position curves that resulted from the immediate free recall of lists of words with different orders of approximation. Introducing the organization into the material presented for free recall gradually changes the

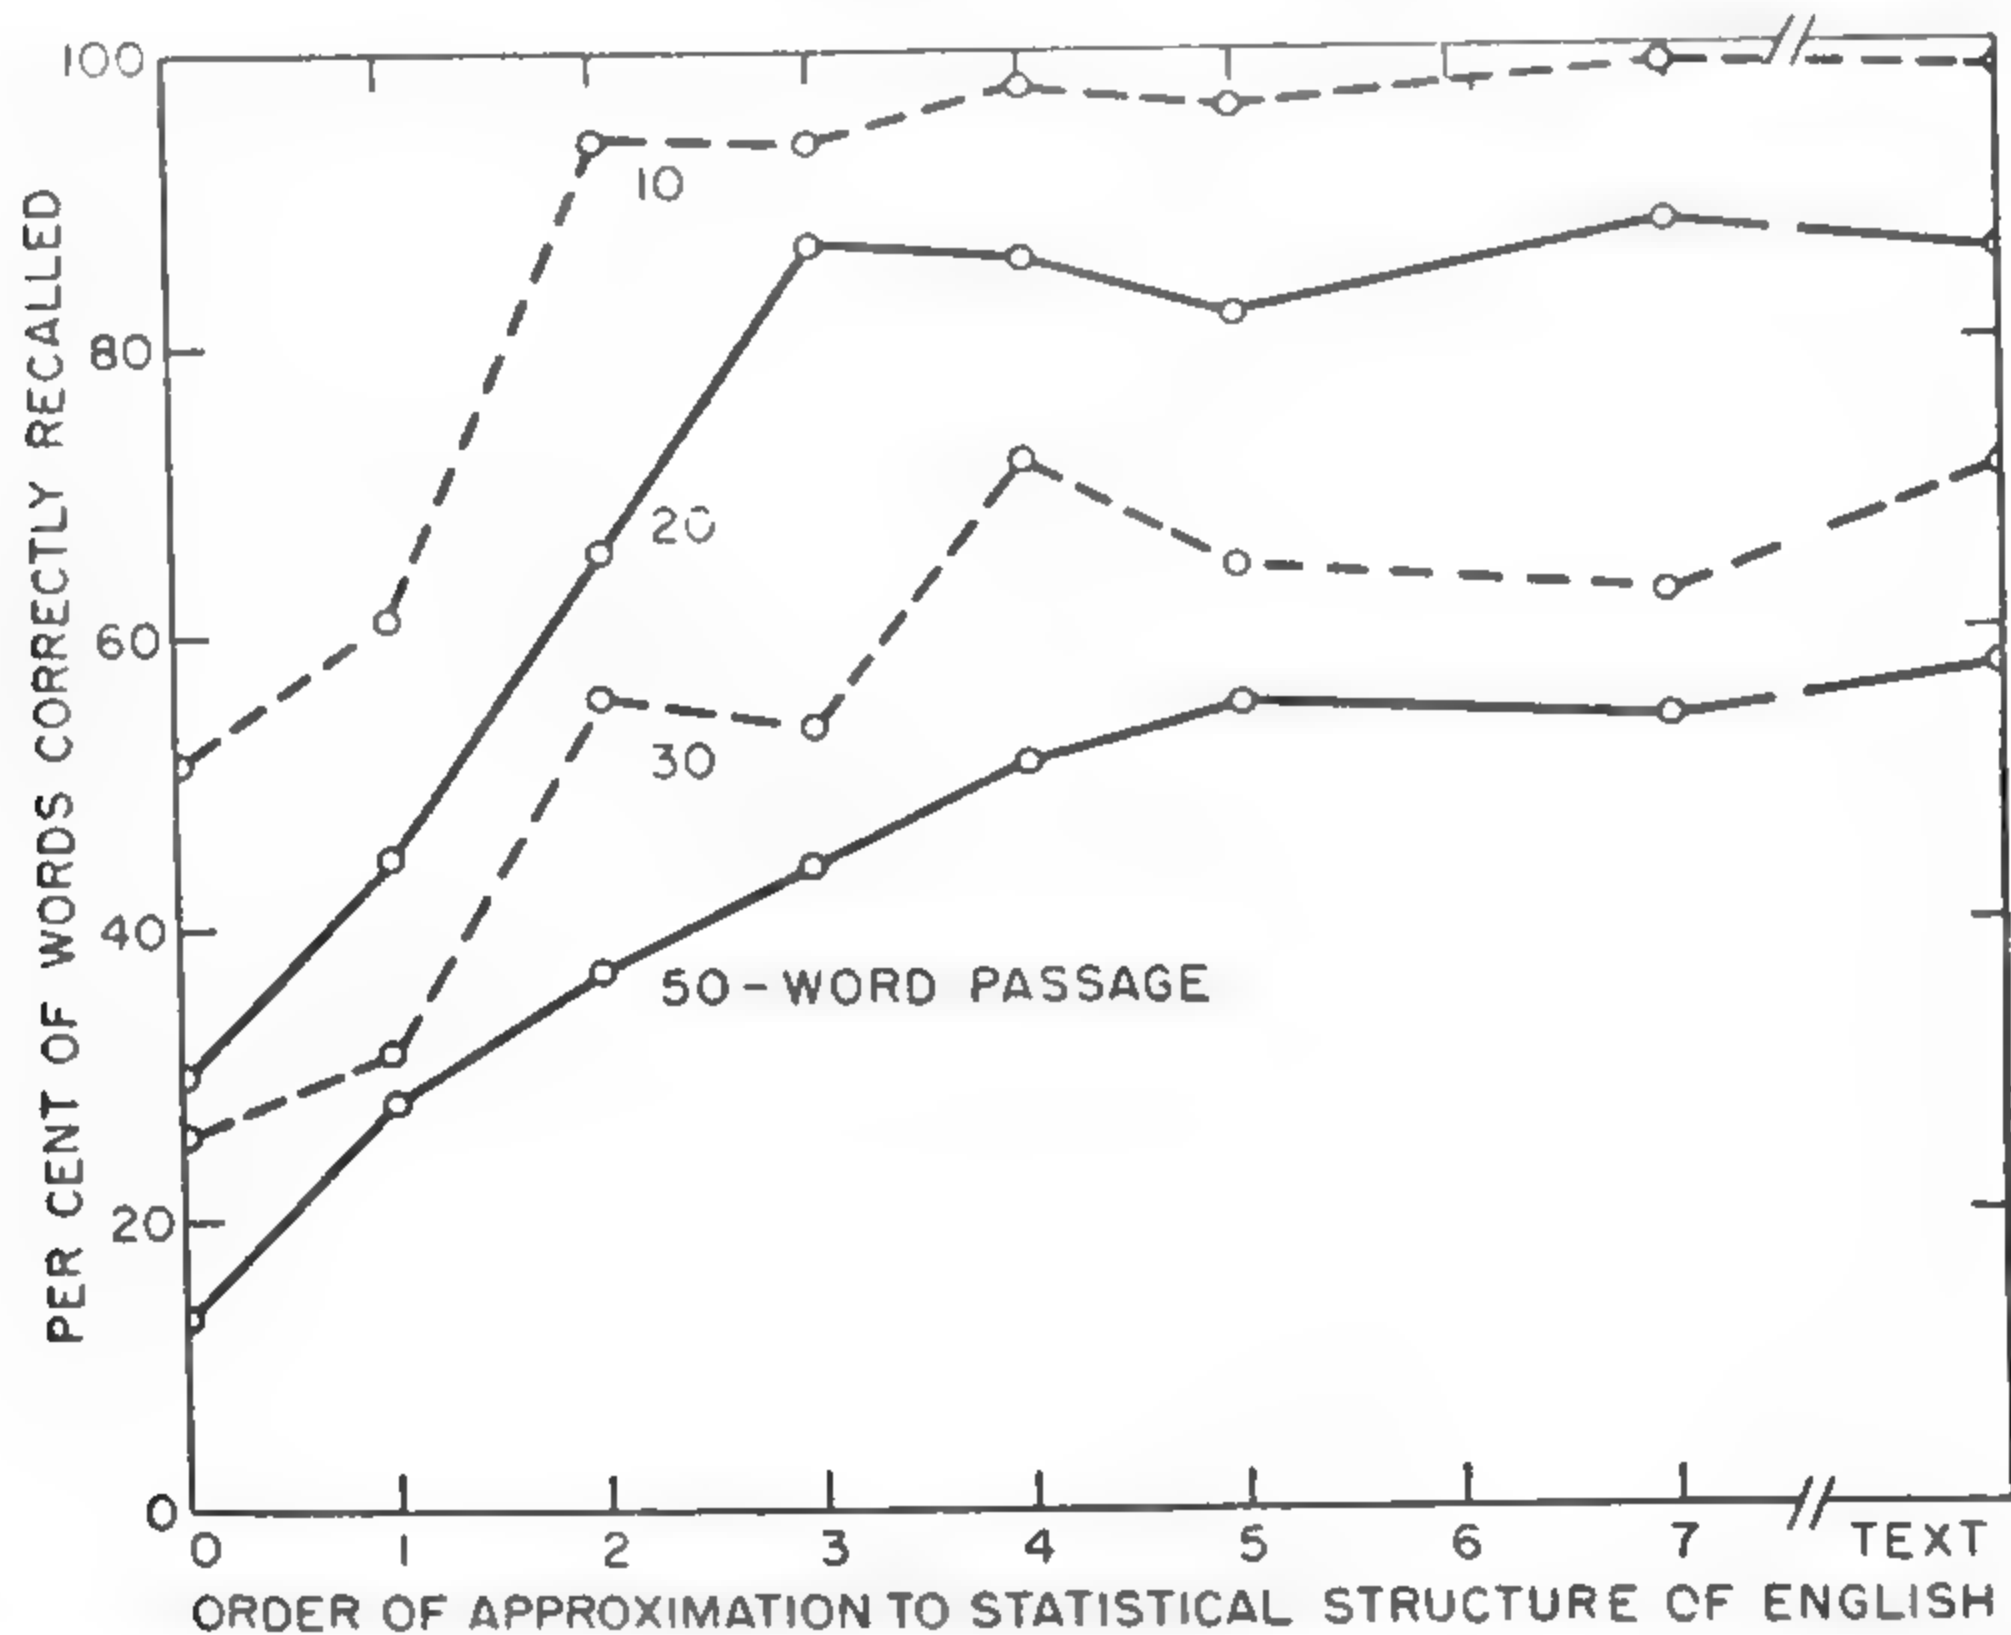


Figure 36. Per cent of words correctly recalled as a function of the order of approximation to the statistical structure of English. (Miller and Selfridge, 1950.)

characteristics of the serial-position curve from those of recall of unorganized material to those of serial anticipation. This seems to be additional evidence that the statistical organization preserves the short-range anticipatory associations in ordinary language. It should be emphasized that statistical dependencies in language do not account for all the greater ease of recall for meaningful material; some of it must be due to the context derived from the content of the material (Marks and Jack, 1952).

Recall and organization in nonsense syllables. In another approach to much the same kind of problem (Aborn and Rubenstein, 1952; Rubenstein and Aborn, 1954) a group of subjects was made familiar with a set of nonsense syllables. These subjects then learned combinations of these non-

sense syllables with various degrees of patterning introduced among the syllables. Under one condition they learned random arrangements of the syllables, under other conditions they learned that there were to be patterns among the syllables and were instructed in the actual nature of these patterns. After learning the nature of the patterns, they memorized the syllables themselves. It was clear that learning the organizational rules beforehand enabled these subjects to remember more individual syllables. However, interestingly enough, they did not profit from the patterns as much as they could have. In terms of *amount of information* (which takes into account the existence of the patterns) they actually remembered less with highly organized than with random patterns. While it is possible that this is merely because they did not have enough practice with the patterns, it suggests that perhaps people, when learning, do not take as much advantage of statistical dependencies as they should.

CHAPTER 8

THE LEARNING OF SKILLS AND CONDITIONS OF PRACTICE

Most of us have at our command an enormous range of skills we have patiently acquired over the years. Some of these skills are highly specialized and some are fairly universal within a given culture. Most American adults, for example, seem to be able to drive an automobile, while only a few of us can operate a machine lathe. The psychologist's interests in these skills are many and varied. Some psychological problems arise, for example, when we ask about the pattern of native abilities that goes into any given skill. Other psychological problems come up when we ask about the interactions of these skills—how an individual may be hindered or helped in learning a new skill by previously acquired skills. In this chapter we shall be concerned primarily with the ways of experimentally analyzing the learning of complex skills and how such learning is influenced by the methods and conditions of practice. We shall largely ignore the question of abilities and defer to later chapters the question of interaction among separate skills.

LEARNING CURVES AND SEQUENCES OF SKILLS

The first thing to occur to anyone given the task of finding out how human beings acquire skills would be to chart the results of practice at the task in question. Psychologists spent a great deal of time doing this in the early experimental studies of the learning of human skills. Often as much as a third of an early textbook on the psychology of learning was devoted to the measurement of improvement with practice and plotting this improvement in learning curves.

There is sometimes much to be gained by this practice. Frequently, in

remedial-reading courses or in special-training courses in reading, the instructor recommends that students keep records of their daily reading rates and plot these on a chart. There seems good reason to suppose that this is a useful teaching device for some skills. Sometimes learning curves are used in industry for monetary incentive scales and for evaluation of the success of particular industrial training programs.

In addition to these more or less practical uses, however, it has sometimes been thought that learning curves would tell us something fundamental about the learning process itself—indeed, perhaps that there was a generalized law about the learning of skills. Unfortunately, such hopes about learning curves have never quite been realized. However, there are general characteristics of some learning curves that are important.

A theoretical derivation of a learning curve. Thurstone (1930a) presented a theoretical derivation of a learning curve based on the assumption that all acts could be classified into successful components of skills and errors which will not lead to successful completion of the act. From this and simple assumptions about rate of error elimination Thurstone derived an equation that describes a learning curve. This curve can represent the acquisition of a remarkably wide class of skills. For example, at least one aspect of the learning of typewriting can be described by this equation, since the results of practice at typing can easily be classified into successes and errors. The derivation, however, would yield no information about a curve describing the effects of practice on speed of typing.

One interesting thing about Thurstone's derivation is that it contains some psychologically meaningful numerical constants. There is one, for example, which concerns the rate of error elimination. If the constant has the value of zero, no errors will be eliminated because the learner will not profit by the performance of an erroneous response, and its probability of reoccurrence will remain unchanged. If, however, this constant has the value of 1.0, each time an error is performed its probability of repetition is reduced to 0. Thus, what is commonly called trial-and-error behavior can be described as happening whenever this error constant is between 0 and 1.0. If this constant becomes greater than 1.0, however, it means that the occurrence of an error eliminates not only this particular erroneous response, but the learner profits enough by this mistake to be able to eliminate possible future occurrences of *other* errors. This, says Thurstone, is the essence of *rational learning*. If a learner can profit by error to the extent of generalizing to other possible errors, it certainly is more than trial-and-error learning.

The possibilities inherent in Thurstone's analysis of learning curves have

never been fully realized. We should note, however, that Thurstone's analysis does not cover all learning curves; it is meant only for those in which the components of performance can be plotted in terms of proportion of correct or erroneous responses. Even here the agreement with data is not as good as with other theories.

Types of learning curves. There are many ways in which components of skills can be measured. For skills in which the components can be classified into errors and correct responses, curves can be plotted which show the way in which errors decrease with practice (trials) or the way in which correct responses increase. In addition we can plot curves which show the length of time it takes to complete a unit of practice at the task. This is the kind of curve one would obtain if the length of time it took an individual to assemble a jigsaw puzzle were measured. Obviously, time per trial decreases with practice, or we do not have evidence for learning. If, for example, the same individual assembled the same jigsaw puzzle several times in succession and there was no decrease in time of assembly, we should conclude that the individual had not shown evidence of learning to assemble the puzzle.

In many tasks we can obtain measures based upon both errors and time. In maze learning, for example, we can plot both the number of blind alleys a subject enters on a given trial and the time that it takes him to go through the maze.

In addition, in some tasks, we may be interested in plotting the way in which accuracy changes with an increase in the number of trials. If we were studying the ability of an individual to learn to fire at a stationary target, for example, we might plot a learning curve in terms of his deviation from the center of the target for each trial. Or, if an individual were tracking a constantly moving target, we might plot the amount of time he stayed on the target for each trial.

It should be noted that the measures of improvement are not all necessarily independent of one another. In maze learning, for example, number of errors and time per trial will be highly correlated. This is simply because the fewer the entrances into blind alleys, the less time it takes to go through the maze. There are fewer mechanical correlations as well. In general, for example, an increase in accuracy will go along with a decrease in the amount of time necessary to accomplish a unit of work. In other words, improvement in one aspect of a task is usually accompanied by improvement in other aspects as well. This is not necessarily so, but it is generally true of learning outside of the laboratory.

In addition there are many other changes in behavior during learning which often do not get charted. Usually there is a change in attitude of the individual toward the task. A task may be regarded as difficult at the outset of learning and as easy after much practice. There may be a change in the degree of tension which an individual shows during learning. In the beginning there may be much tension and anxiety which may be reduced by the learning of a task. Finally the mode of attack of the learner will most surely vary throughout the course of learning. Thus it is clear that there are many aspects of the change in performance with practice, and some of these are often ignored.

Many times learning curves are obtained by averaging together the performance of a number of individuals at each trial. The learning curve then is a kind of average learning curve. "Kind of average" is sometimes apt, for there are frequently serious distortions introduced by this procedure. This has been pointed out many times, most recently by Estes (1956). Most generally this is true because the form of the curve for averaged measurements is not necessarily the same as that for the individual subjects. For example, suppose we plotted learning curves for a group of individuals who always learned instantaneously (as would be the case if the learners had "insight"). Thus on some one trial an individual's score would go from zero to the maximum possible score. If some of these individuals learned at *different* trials, the result could be a smooth, gradual learning curve for the average of all these individuals. If we looked only at the average curve we should infer that the learning of this task proceeded in a gradual way, with, perhaps, more learning in the early trials than in the later trials.¹ Thus, because learning curves for individuals are not always the same as those for groups, we must be careful about the inferences we draw.

Plateaus. If we look at learning curves for individuals an interesting phenomenon occasionally turns up. This is illustrated by the curve in Figure 37. The task this individual was practicing was receiving telegraphic code. Notice that there was a period of about six weeks during which the subject did not appreciably increase his ability to receive code signals. After this period, there was a sudden spurt in letters received per minute which gradually leveled off to a final asymptote. Such a period of little or no improvement is known as a plateau.

The particular curve in Figure 37 comes from a well-known older study by Bryan and Harter (1897, 1899), who were trying to determine the reason for this plateau, which they thought fairly typical of learning tasks

¹ This would be true if the distribution of occurrences of "insight" were logarithmic.

like receiving code. They put forward the notion that a plateau occurs when an individual has learned all of the primary skills in a complicated task but has not yet organized them into broader units so that he can learn the higher-order skills. They argued that learning any skill like telegraphy involves a hierarchy of habits; that is to say, the learning of the higher-order skills depends upon mastery of the lower-order skills. In learning to receive telegraphic code, for example, individuals first learn to receive individual letters. Then as they master the skill involved in this, they can take advantage of the redundancy, or predictability, in English and learn to receive by words instead of by letters. Finally, they can receive phrases or words as units.

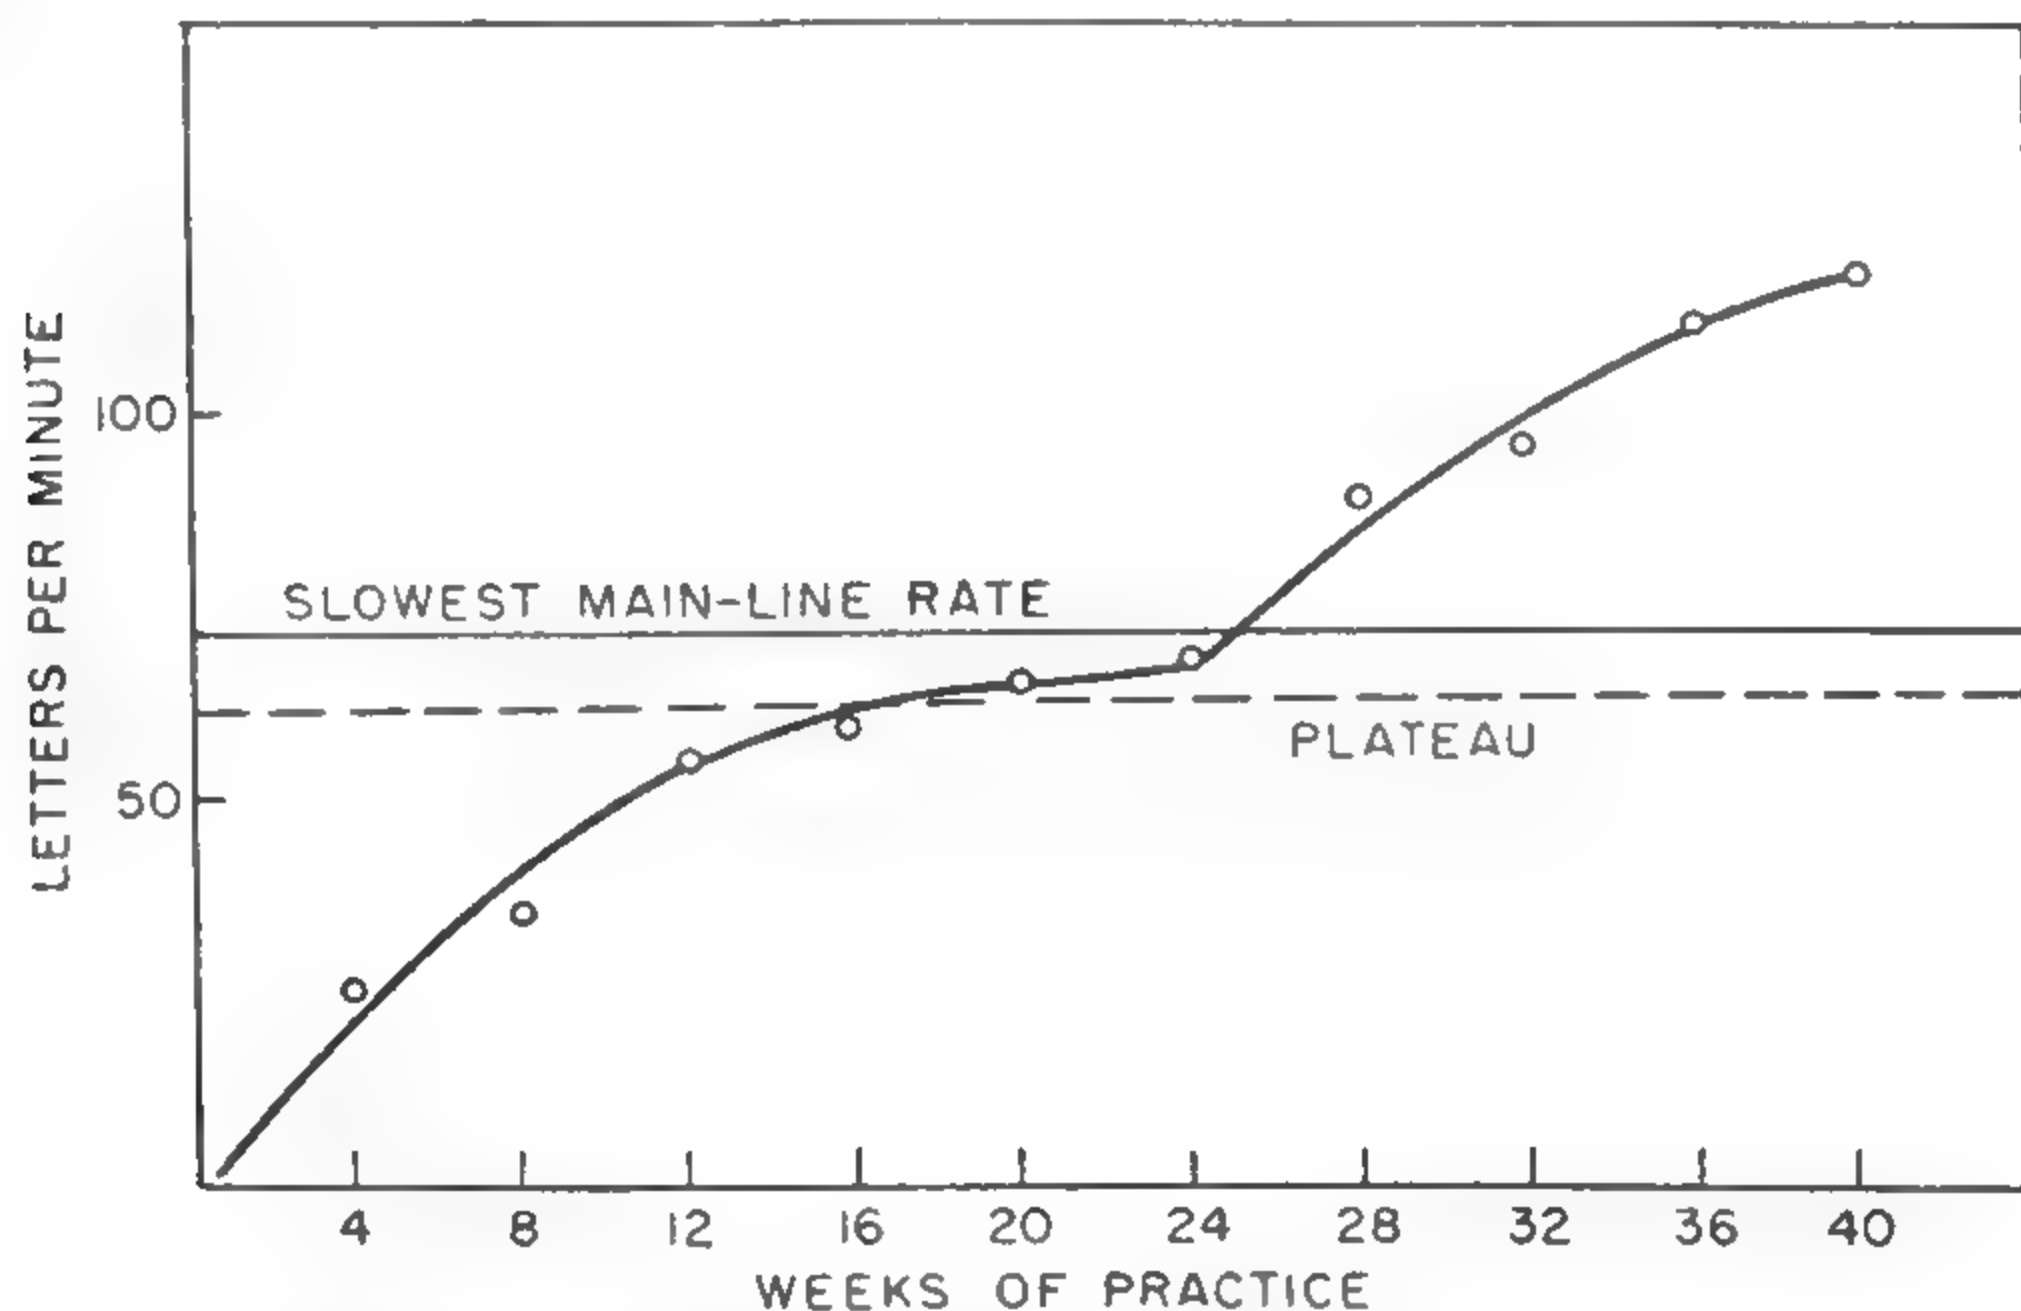


Figure 37. Learning curve for receiving telegraphic code. A plateau occurs in the middle of the curve just before the minimum rate is reached. (After Bryan and Harter, 1897.)

Bryan and Harter argued that the plateau ought to occur at the transition between learning to receive individual letters and learning to receive words. They were able to justify this interpretation by a study of the way in which people learned to receive jumbled letters rather than real words. They found that the plateau for meaningful material occurred at just about the point at which there was no further improvement in the ability to receive disconnected letters and words. Thus, the plateau occurs where the limit on the lower-order habits has been reached and where the higher-order habits have not yet begun to appear. The results of this comparison of disconnected and connected material can be seen in Figure 38.

It is obvious that the notion of hierarchy of habits is useful whether we are interested in plateaus or not. Something like this hierarchical structure

must occur in all tasks in which there is opportunity to recode the material into larger and larger units. Thus, learning to typewrite, read music, and the basic skill of learning to read itself must involve something like hierarchies of habits.

Plateaus do not always occur. Indeed, in telegraphy they seem to be the exception rather than the rule (Taylor, 1943). The transition between orders of skills is more often a smooth one. Also, there seem to be other causes

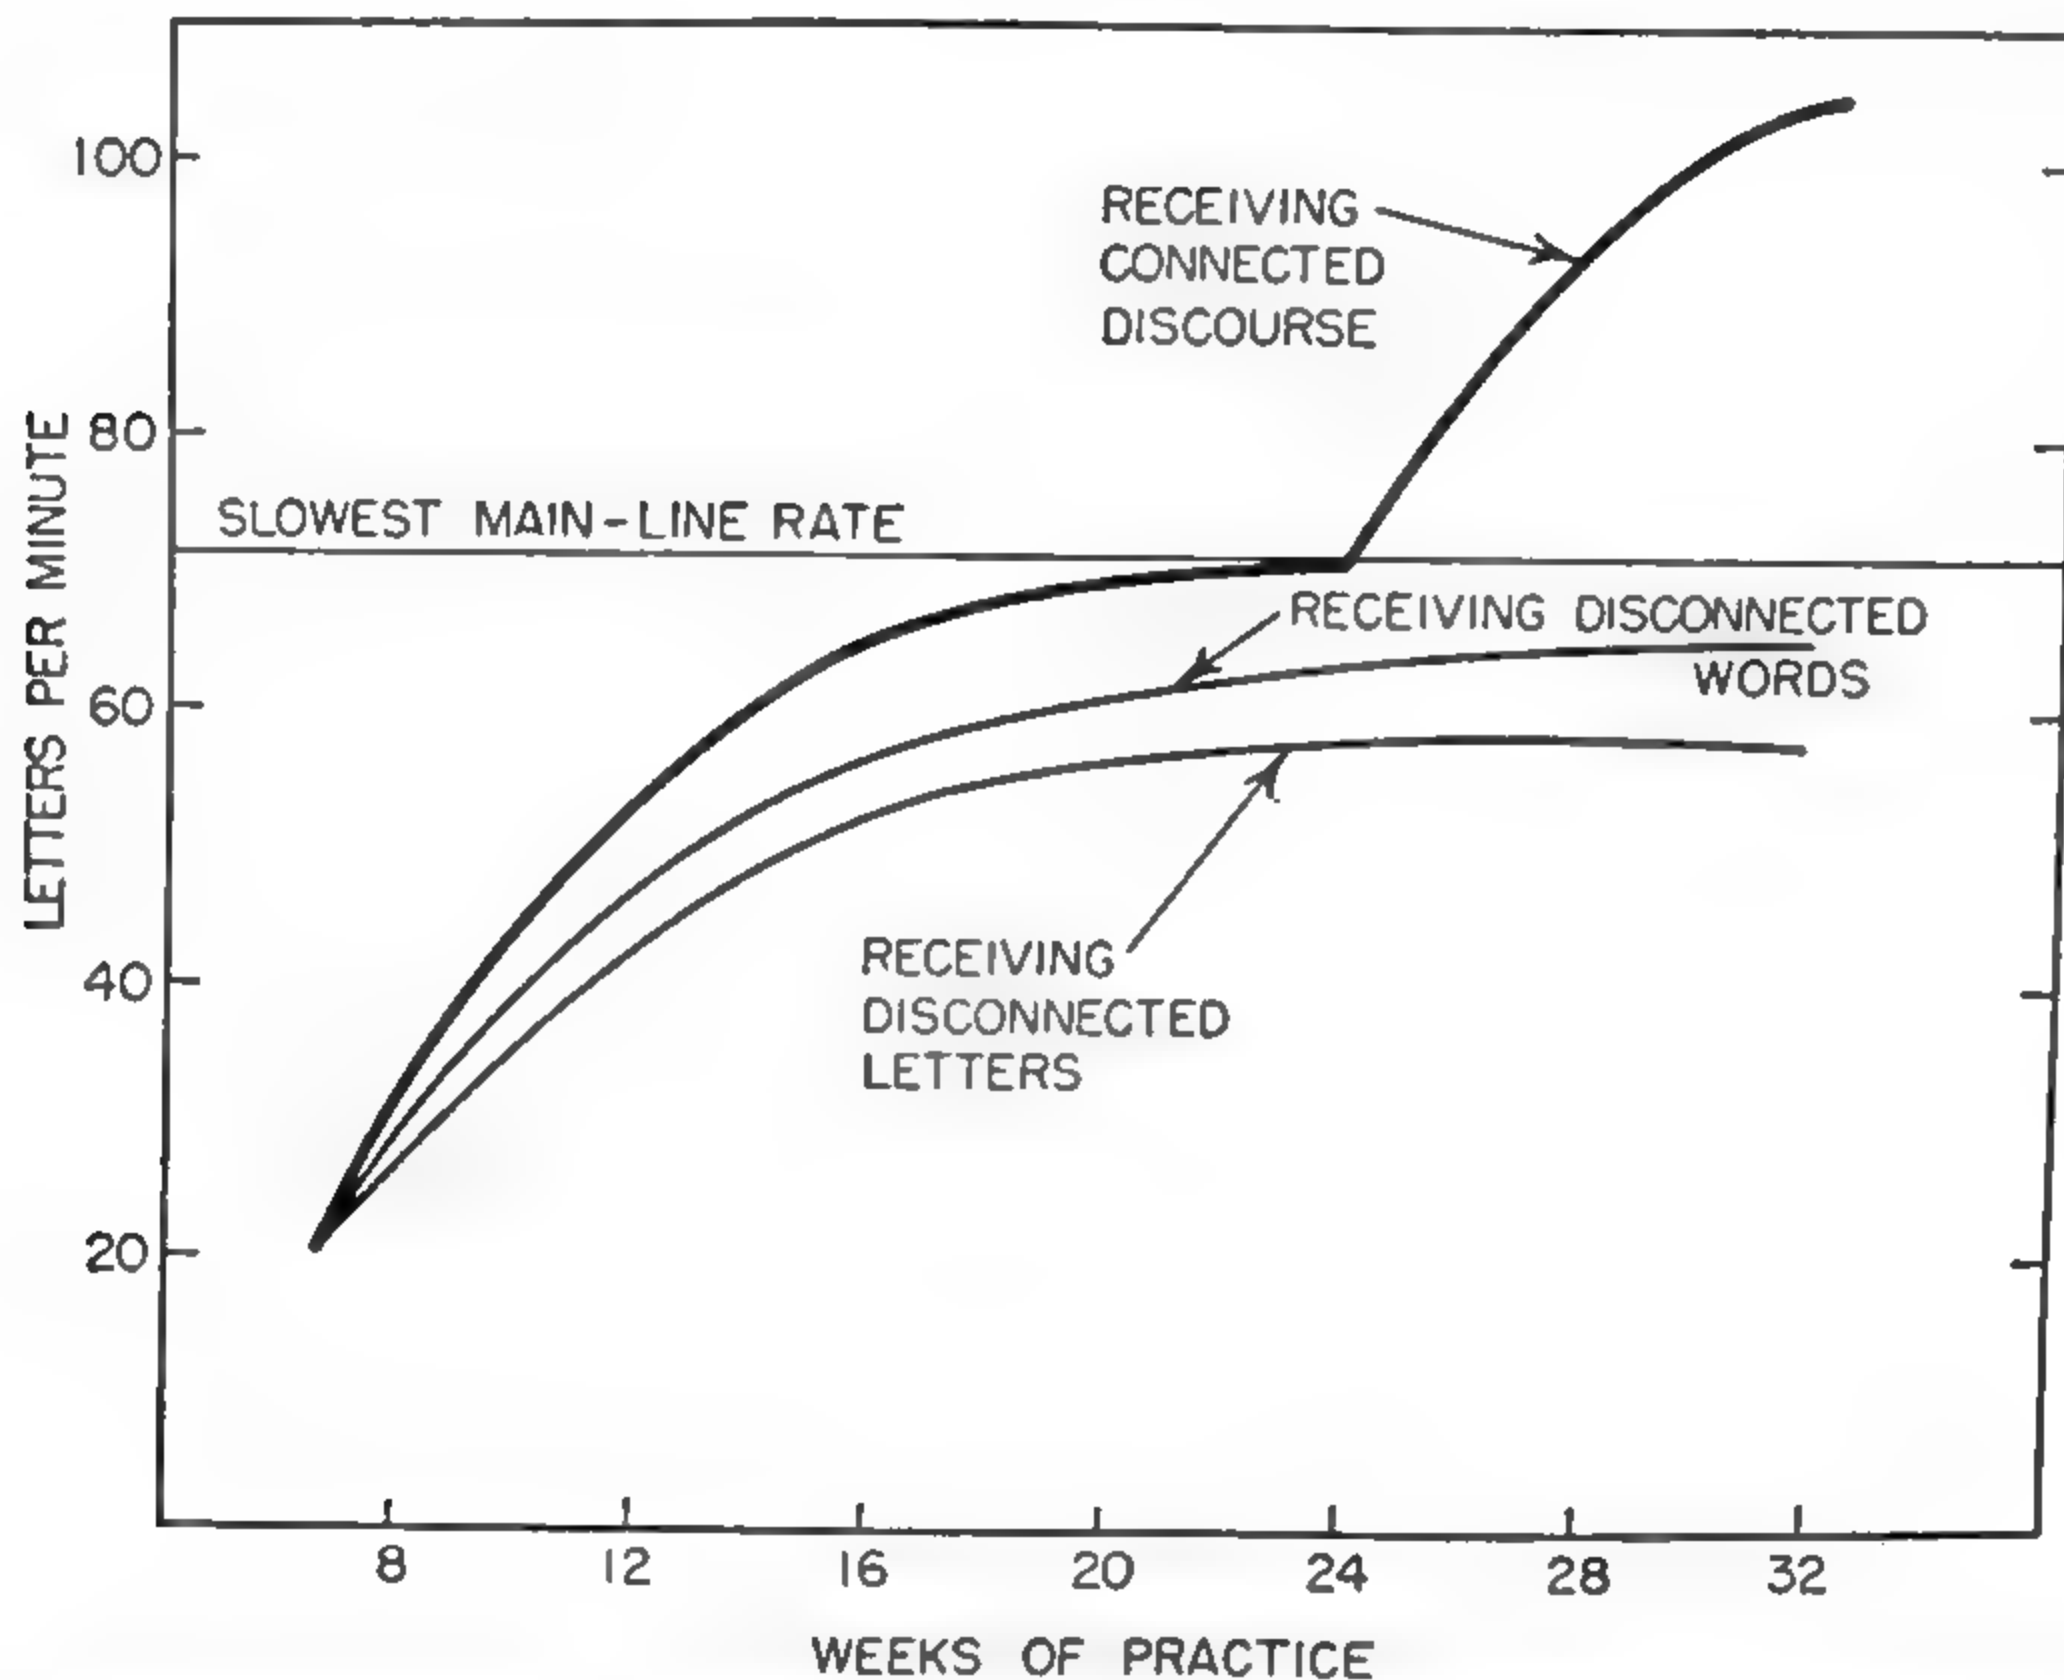


Figure 38. Learning curves for receiving connected discourse, disconnected words, and disconnected letters. Notice that the plateau in the curve for connected discourse occurs at the point where there is no further improvement in the curves for disconnected material. (After Bryan and Harter, 1899.)

of plateaus in the learning curve. Swift (1918) argued that plateaus are caused by a decline in the learner's motivation. He pointed out that plateaus for all learners occur in about the same place; furthermore they occur in many tasks which do not involve a hierarchy of habits. Book (1925), in his handbook on learning to typewrite, located the lapses in attention and drop in motivation at the point where letter habits were making a transition into word habits. Thus the failure to improve generated discouragement, which furthered the failure to improve. Individuals may become so discouraged at failure to improve that they lose motivation for further practice. Indeed, plateaus may well be a fact which keeps adult learners from ac-

quiring such skills as playing tennis or playing the piano. A well-motivated adult can learn the fundamentals of these skills quite easily, but there then follows a long period when devoted practice brings no fruits. At this point interest drops and the individual will cease taking lessons.

Frequently, the limit of performance which learners finally reach is a plateau. The author doubts that his typing has improved much in the past fifteen years, but there is certainly room for improvement. A little practice would eliminate some bad habits and bring in some new levels of organization which would no doubt quickly raise his level of typing. Thus, while he is currently at a limit of performance in typing, practice would raise that limit. There are surely real physiological limits to our ability to perform specific skills; we cannot type faster than our fingers will move. Most of us, however, seldom push ourselves to this limit.

CONDITIONS OF PRACTICE

The learning of any particular skill is a function of the conditions under which it is practiced. The shape of the over-all learning curve may be altered by change in conditions, or the presence or absence of particular components of the skill may be altered. In this section we shall examine the influence of some of the more important conditions of practice which have been experimentally studied. We shall restrict ourselves to conditions that apply to many tasks, since we are interested in general principles of learning rather than the characteristics of certain tasks. To take an example: We might be interested in the influence of the location of individual letters on the typewriter keyboard, since this could tell us something about general problems of habit interference, but we should not be interested in the influence, say, of depth of keystroke on learning to typewrite. The latter problem may well be of interest to those concerned with the specific skill of typewriting, but it is of little general interest in the psychology of learning.

The Distribution of Practice

The condition of practice that has been explored experimentally more extensively than any other is that of the distribution of practice through time. For a wide variety of tasks, the question has been asked, Is it better to practice a task with as little interruption for rest as possible, or is rest beneficial to learning and performance? For most of the experimental tasks that have been studied, the results show that some rest is better than none at all, and for some tasks the advantage in increased skill is quite large.

The problem is complicated; there are many tasks for which distribution of practice has little effect, and there are a few for which rest periods actually seem to decrease ability. Because of the complexity of the problem and because of the theoretical interest in the distribution of practice, it will be necessary for us to examine this condition in considerable detail. Let us first look at the experimental findings and then at some of the theoretical questions.

Experiments on the Distribution of Practice

A simple experiment. A well-known experiment by Lorge (1930) illustrates the basic phenomenon in the distribution of practice. Lorge studied the effects of continuous practice for 20 trials compared with the effects of practice in which a rest period of one minute or a rest period of one day intervened between each trial. The tasks he used were mirror drawing, mirror reading, and code substitution. He found that both cases of distributed practice resulted in better learning than massed practice. Figure 39 shows the results for the mirror-drawing task. We can see that the time per trial for the two groups with distributed practice dropped much more rapidly than the curve for the massed-practice group; the difference is apparent almost from the very beginning of practice. Notice also that the difference

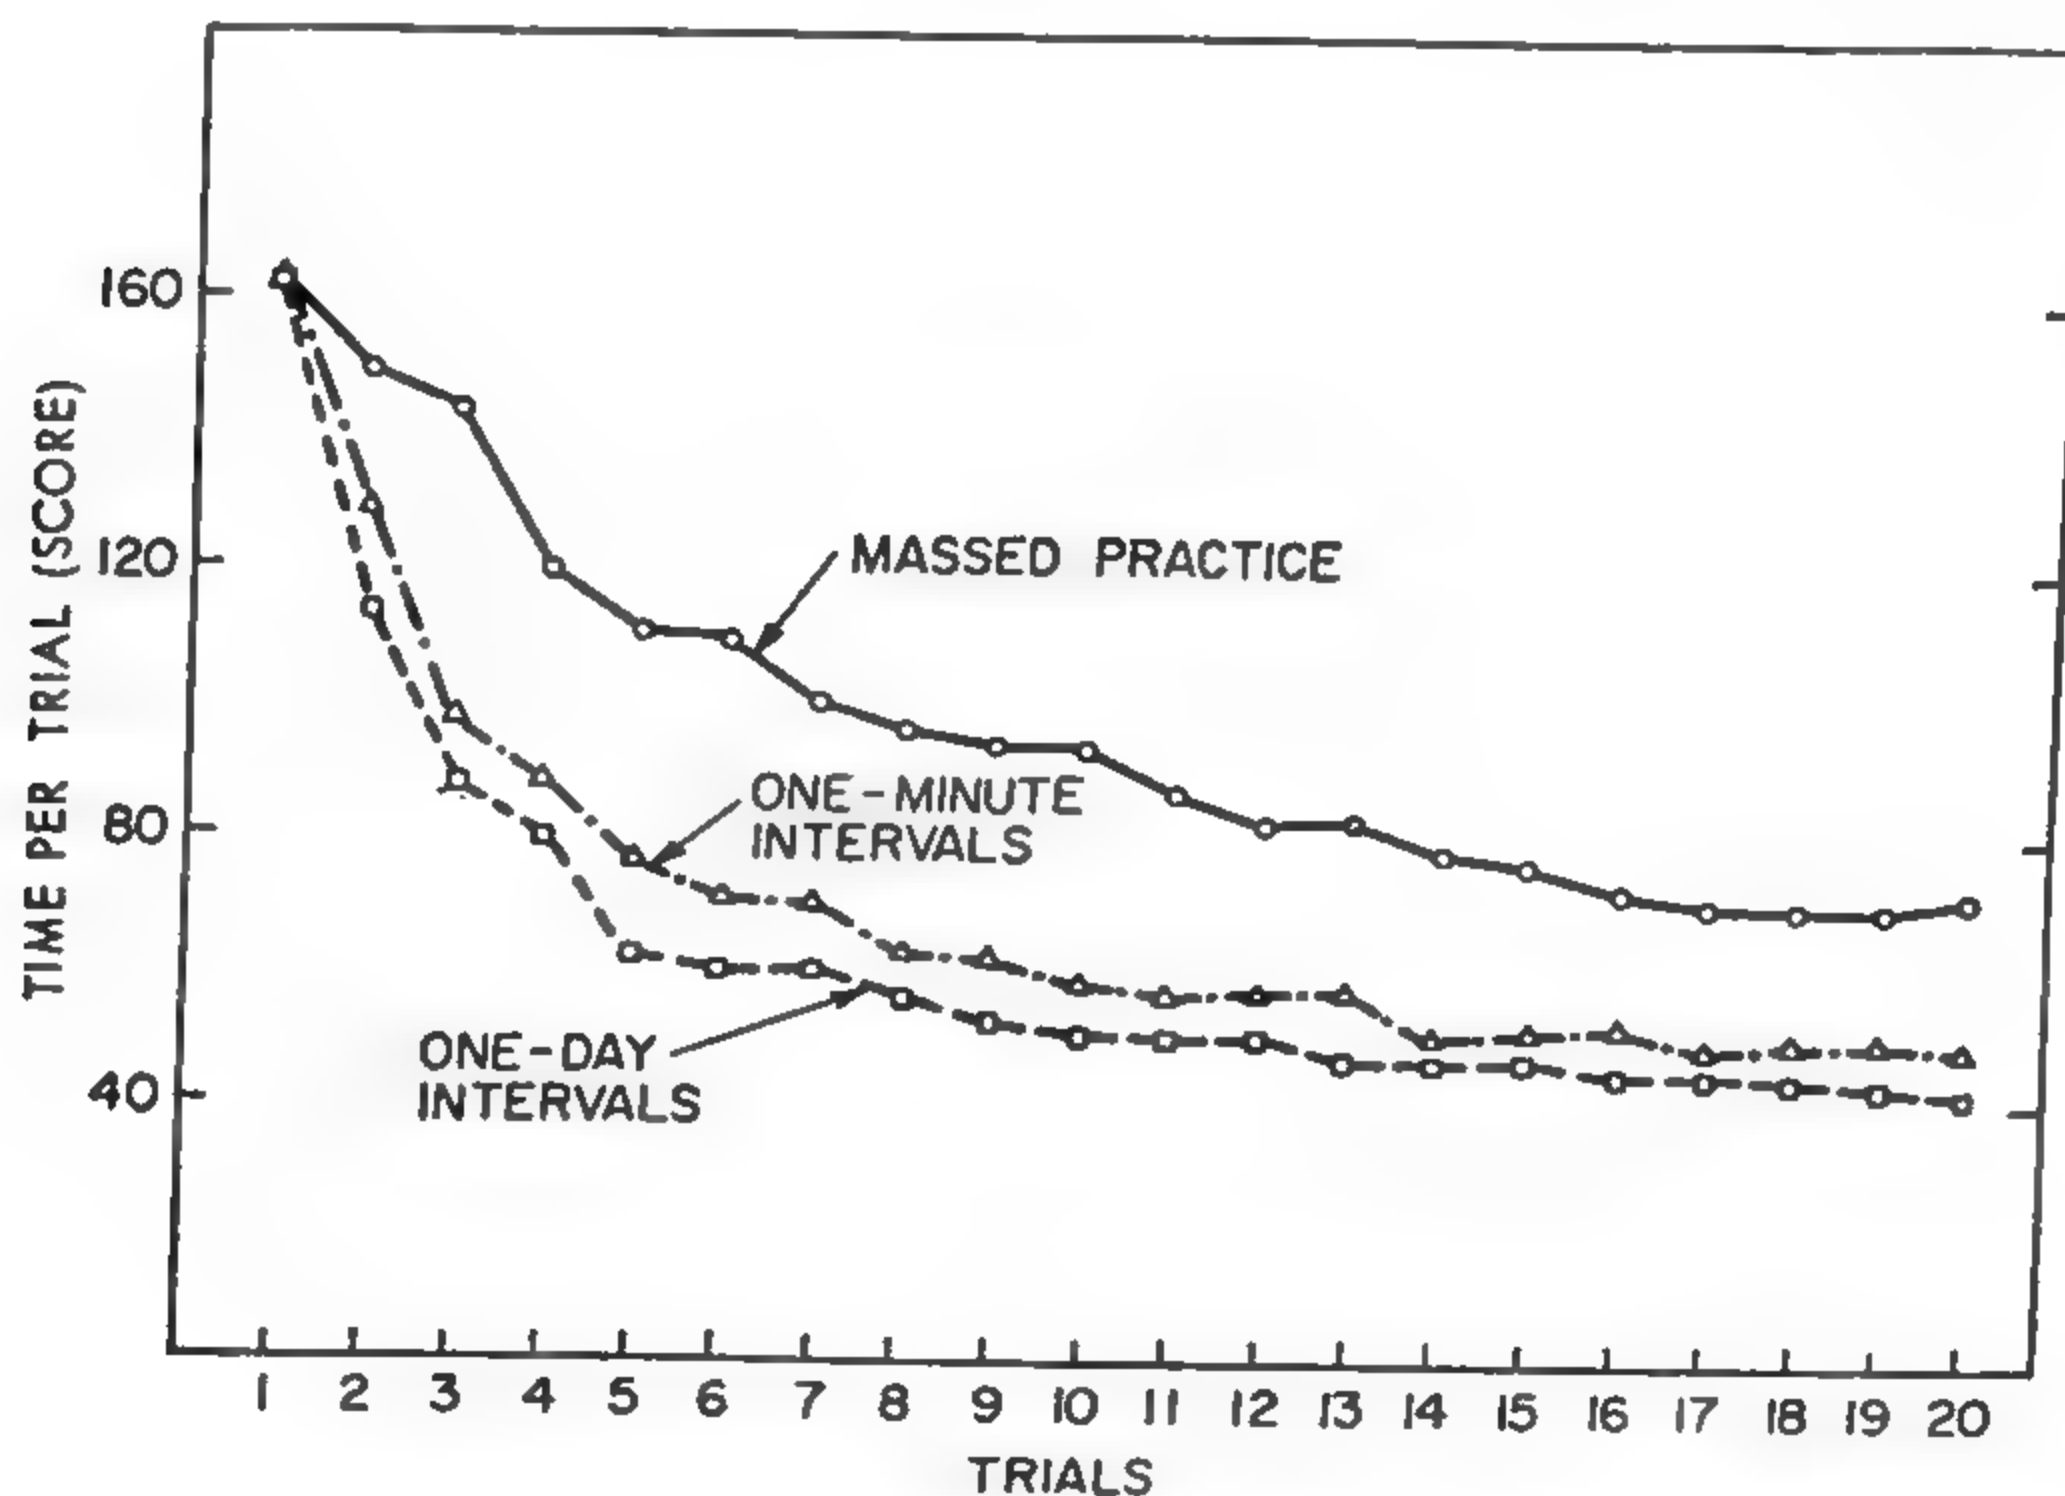


Figure 39. The effect of distribution of practice upon mirror drawing. Notice that, apart from the very first trials, there is almost a constant difference between the different groups. (After Lorge, 1930.)

between the group that rested for one minute between each trial and the group that rested for one day is very small.

This experiment illustrates the basic effect of distribution in a perceptual-motor task of moderate difficulty. As is typically the case for such tasks, the difference produced by some distributed practice compared with none is large and consistent.

The effects of distribution when practice is unequal. In the experiment just cited the amount of practice for the various groups was held constant and the time over which practice occurred was varied. What happens, however, when the time over which practice occurs is held constant and the amount of practice varied? In this case, a group of subjects under massed practice would receive more practice than one under distributed practice.

Duncan (1951) studied the effect of distribution of unequal practice on the acquisition of the ability to keep a small stylus on a constantly moving target (rotary-pursuit test). He allowed some of his subjects to practice continuously; the remainder of his subjects were periodically interrupted for a rest period. The rest periods actually occupied two-thirds of the practice period, so that the group practicing continuously had three times as much practice. Despite this heavy balance in favor of the massed-practice group, distributed practice produced better performance. The distributed practice group maintained the superiority it had shown during the initial learning after a pause of 10 minutes. Therefore it is clear that the effects of distribution are powerful and consistent.

Length of work and rest periods. It is reasonable to suppose that the relative effectiveness of distributed practice would depend upon the absolute and relative lengths of the work and rest periods. There are a number of experiments that explore this problem and provide us with important clues about the underlying mechanism responsible for the distributed-practice effect.

Kientzle (1946) studied the effect of varying the rest period while holding the work period constant. The measure of performance she used was the number of alphabet letters subjects could print upside down in a one minute trial. She varied the rest intervals between trials from zero to seven days. Her results can be seen in Figure 40. In general, her experiment showed that for this task small rest periods resulted in a great improvement over no rest between trials, but that beyond 45 seconds, not much advantage was gained by increasing the rest interval.

In another experiment on the alphabet-printing task, Kimble (1949b) showed with trials of 30 seconds duration that the advantage of spacing

trials kept on increasing from zero to 30 seconds rest. Since Kimble did not study distributions greater than 30 seconds, we do not know whether there would have been any advantage gained from even longer rest periods.

In the two experiments just discussed, the duration of the work period was kept constant. What is the interaction between the length of the work period and the length of the rest period? In one experiment (Kimble and Bilodeau, 1949) the task was rapid manipulation of small pegs, and it was found that shortening the work period is much more important than length-

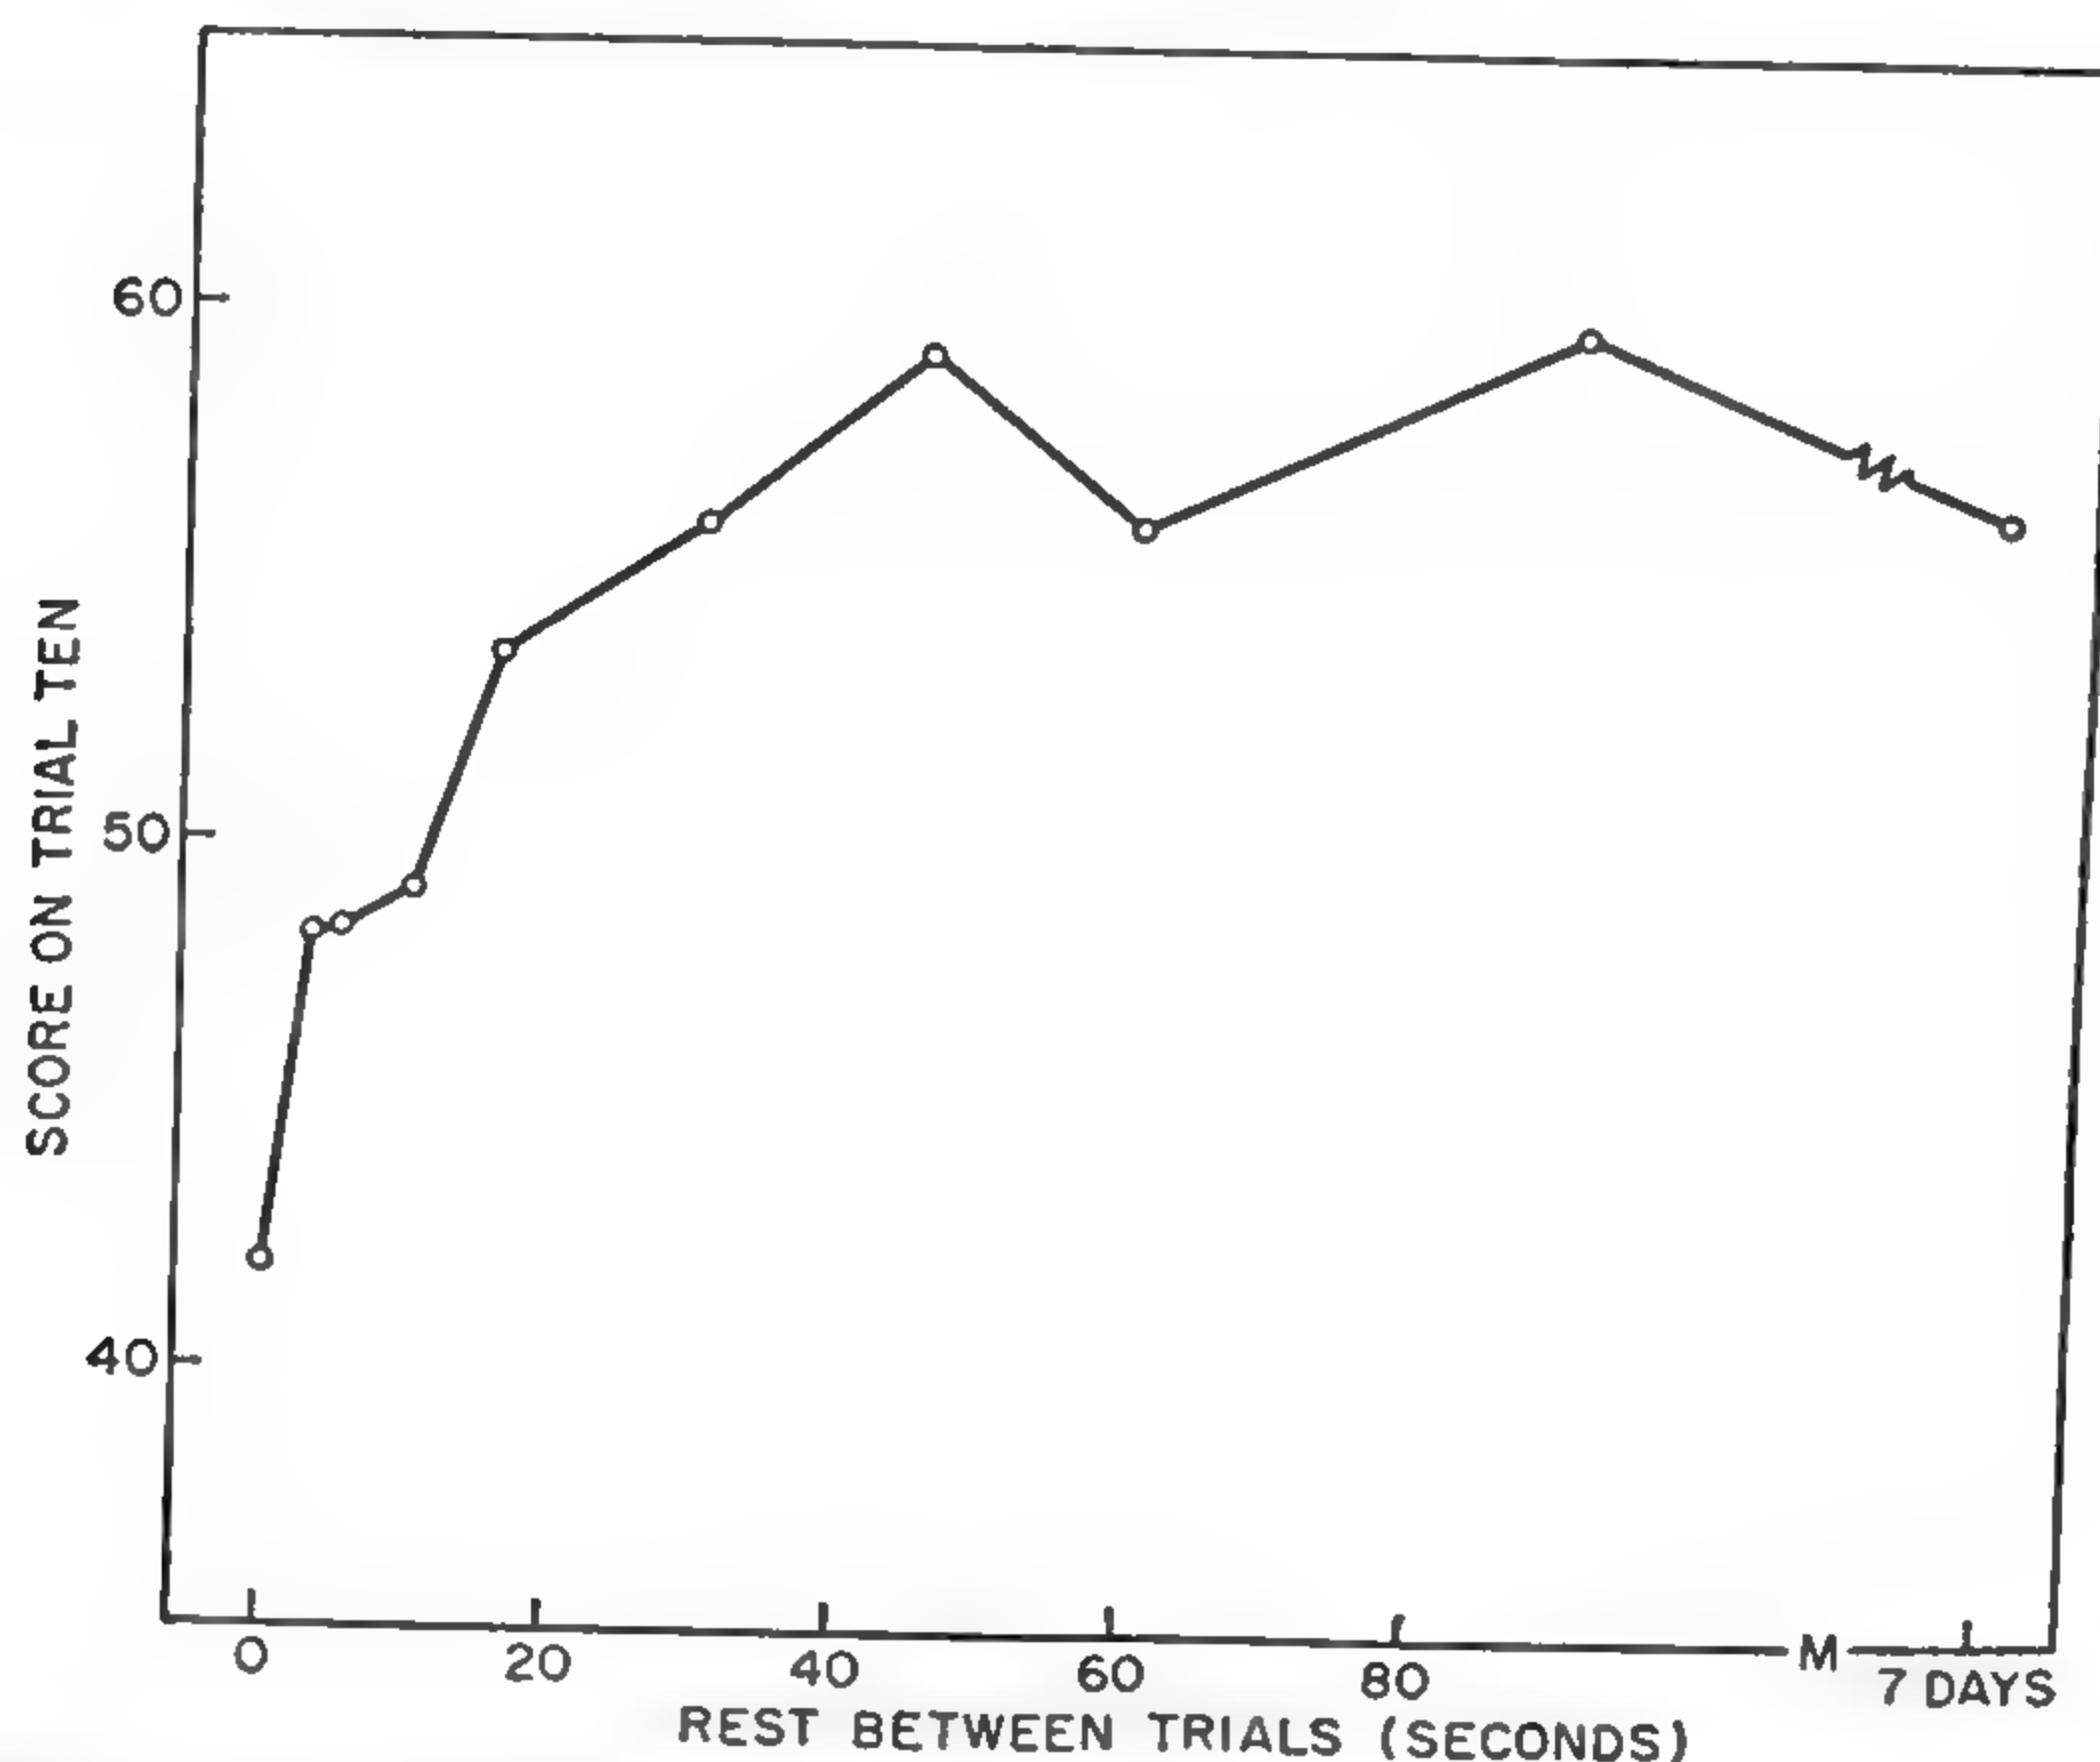


Figure 40. The effect of time between trials on performance of a learned act on a particular trial. Notice that for this task (printing upside down) rest intervals greater than 45 seconds do not appreciably increase level of performance. (Data from Kientzle, 1946.)

ening the rest period. Thus, the optimal rest period for any given task probably depends upon the duration of the work period during practice. From the available evidence, it is likely that short work periods and short rest periods are most beneficial to practice, but the actual duration of the work period will clearly depend upon the nature of the task. It will not do, for example, to break practice into unnaturally short units for the task in question; interrupting someone solving a puzzle before he has reached solution may well have disrupting effects upon practice.

The interaction of distribution and nature of the task. A large literature, only a fraction of which is cited here, supports the contention that some dis-

tribution of practice is better than none for nearly all psychomotor skills. We do not know very much about how distribution of practice interacts with the characteristics of specific tasks, however. For theoretical reasons (see below) it has commonly been supposed that distribution ought to be more effective for tasks involving a large amount of physical work and effort expenditure than for those involving little work and effort (see Ellis, 1953), but there has been little systematic exploration of task-distribution interactions for psychomotor skills.

Underwood and his associates, however, have studied rather exhaustively the relationships between distribution of practice and various conditions prevailing in rote verbal learning. The results are rather complicated. For example, one study (Underwood and Goad, 1951) showed that distributed practice was superior to massed practice for learning lists of adjectives *by the method of serial anticipation* if the similarity between adjectives was high. If the similarity between adjectives was low, massed practice was as effective as distributed practice. On the other hand, irrespective of the similarity of items, learning adjectives *by the method of paired associates* was unaffected by distribution (Underwood, 1951). Again, to make the picture more complicated, Oseas and Underwood (1952) find that distributed practice benefits the learning and retention of concepts, when the method requires subjects to abstract concepts by being exposed to successive examples of the concept. Furthermore, comparisons of the results achieved by Underwood and his associates with those achieved by another study (Hovland, 1949) suggest that the level of skill or sophistication of subjects in laboratory learning may interact with the distribution of practice (Underwood, 1953b).

The net result of all the intensive work on the effects of distribution of practice on a particular kind of learning has been to show that it is extremely difficult to apply a simple theory to all of the known phenomena of the distribution of practice. This is important, because much of the interest of experimental psychologists in the distribution of practice is not so much in the effect itself as it is in the means the effect has provided for testing certain theoretical notions about learning and skilled behavior. We shall return to some of these problems when we consider theories of the effects of distributed practice.

The permanence of the effects of different distributions. We have seen that for the vast majority of psychomotor skills, distributed practice is superior to massed practice. For many verbal tasks this is also true. It is to be noted that wherever the effects of distributed practice have been studied

in realistic training programs, distributed practice has proved to be superior to massed practice. For example (Crawford, et al., 1947), in one study of training fighter pilots to fire at towed targets, the results clearly showed that if a fixed amount of ammunition was distributed over many training missions the percentage of target hits was as much as five times better than if all the ammunition was fired on one mission.

An important question we have not yet answered concerns the permanence of the effects of distributed practice. If we teach two different groups of people a simple task, one group practicing with distribution and the other with massed practice, the chances are that we shall establish a difference in performance at the end of practice. If we wait a period of time, however, before measuring retention, will the difference between these two groups disappear or be maintained?

One kind of experiment in which the permanent effects of distribution have been important is found in studies of *reminiscence*. Reminiscence is a somewhat curious name, only remotely related to the usual context of the word, but it is commonly used in experimental research on learning, and we shall use it here. Reminiscence occurs when individuals practice at some skill, stop, and after an interval of time are tested for further performance. The characteristic defining reminiscence is that there is an actual improvement instead of a decrement in performance (as there would be if forgetting took place).

The original studies of reminiscence occurred in rote verbal learning and led psychologists to believe that it was a fairly general phenomenon in such learning. More recent studies suggest that in rote verbal learning it is limited to the learning of material like nonsense syllables when practice is severely massed (Hovland and Kurtz, 1951). Under these conditions there appears to be a slight temporary increase in the ability to recall nonsense syllables after a little time has elapsed.

If the phenomenon of reminiscence were limited to rote verbal learning of nonsense syllables it would not be of much significance, since it occurs in limited amounts and only under highly specialized conditions for even this kind of learning. It turns out, however, that considerable amounts of reminiscence occur in the learning of psychomotor skills.

The importance of reminiscence to the problem of the distribution of practice is that it seems to be a function of the massing of practice. In the case of nonsense syllables, no reminiscence occurs with distributed practice, but some does occur with severely massed practice. The implication is clear; with the passage of time some of the difference induced by massed

and distributed practice disappears—the massed-practice group performs more like the distributed-practice group. Two important questions arise. Does this difference in reminiscence between massed and distributed practice occur in psychomotor skills? If so, does the difference between massed and distributed practice entirely disappear after a time interval?

The answer to the first question appears to be clear enough. In psychomotor skills, the amount of recovery during a rest period seems to depend upon the amount of massed practice that has gone before. The second question poses a more complicated problem. Under some conditions the difference between massed and distributed practice does entirely disappear. For example, in an experiment by Kimble (1950) on a psychomotor skill, subjects practiced at the task under either severe massed or distributed practice. After a six-minute rest period, despite large initial differences in performance between the two groups, the group that practiced under massed conditions did as well as the one trained under distributed conditions. In other words, reminiscence from massed practice made recovery from the effects of massed practice very nearly complete.

This case of complete or nearly complete recovery (some ambiguity exists because of warming up when returning to a task after rest) is by no means the rule. It may, indeed, be the exception. For much the same kind of task, Jahnke and Duncan (1956) discovered that a difference produced by massed or distributed practice remains after an interval as long as four weeks, although the difference between massed and distributed practice after so long an interval is smaller than it was originally. Reminiscence occurs, but not so that the difference between the massed and distributed groups is wiped out.

In verbal learning, much the same sort of situation occurs. Sometimes reminiscence shown by a massed-practice group of subjects will wipe out performance differences between massed and distributed practice and sometimes it will not. Even if reminiscence as usually defined does not occur, it is possible that the difference between massed and distributed practice could disappear in a later recall. This could happen if both the massed and distributed group forgot enough items to obtain equally poor recall scores. Again, on this issue, the results of different experiments are not in agreement. Retention of verbal material learned by rote methods (either serial-anticipation or paired-associate) may or may not be enhanced by distributed practice (see Underwood, 1953a). The reason for this lack of uniformity in experimental findings is important, and we shall return to it in a discussion of the theoretical issues in the distribution of practice.

Transfer of the effects of massed practice. It is clear that massed practice has some detrimental effects on performance, and these effects may or may not be permanent in nature. They turn up more consistently and clearly in psychomotor skills than in rote verbal learning. One important question has concerned the physical locus of these effects, for if distribution is more effective in behavior which clearly involves large scale muscular movements, the suggestion immediately comes to mind that this might be because of some fatiguelike inhibition in the peripheral response mechanisms. Consequently a series of studies have been performed with a view to discovering the locus of the effects of massed practice. These experiments have usually been on bilateral transfer.

In bilateral transfer, a subject is taught to perform some simple act of coordination by practicing with one of his hands only. After he has reached a certain level of skill, he is tested for this act by the other hand. If the subject is able to do better with the unpracticed hand than a completely unpracticed subject would do, we say that bilateral transfer has taken place.

Let us look at an example of the bilateral transfer of reminiscence (Irion and Gustafson, 1952). In this experiment subjects practiced continuously for five minutes at following a moving target (pursuit rotor). After practice, one group of subjects immediately received five minutes additional practice with the other hand, while subjects in a second group rested for five minutes before changing hands. It turned out that the group that rested was significantly superior with the other hand to the group that transferred immediately. Thus the *detrimental* effects of massed practice transferred from one hand to the other; they were not limited in location to the muscles involved in the original practice. This effect has been repeatedly confirmed (see Rockway, 1953), and it must be seriously considered in the theoretical issues concerning the distribution effects.

Theories of distributed practice. Current attempts to account for the comparative effects of massed and distributed practice for the most part have emphasized the role of inhibition from massed practice. Some people have argued that the advantage from distributed practice is mainly the result of rehearsal during the rest intervals, but the evidence is against this notion on a number of counts. First of all, many experiments have been controlled for possible rehearsal, and these have still found superiority in distributed practice. Secondly the effect is more stable and larger in magnitude with psychomotor skills than with verbal material, and, intuitively, it would seem that verbal material would be easier to rehearse away from the testing situation than psychomotor skills. Finally, the effect also occurs in animal

learning (Bunch and Magsdick, 1933; Holland, 1953), where rehearsal is improbable.

We have already had occasion to examine an important theory about inhibition. In the chapter on extinction, we examined a theory of extinction that attempted to account for the experimental evidence on the basis of two inhibitory states. One inhibitory state, reactive inhibition, is the simple result of responding. It builds up through continued responding and disappears through rest. A second inhibitory state, conditioned inhibition, depends upon the first; the organism learns not to respond (or, perhaps, not to respond well) as the result of being induced not to respond through reactive inhibition. This notion, originally advanced by Hull (1943), has been modified and extended by several investigators. It was originally put forward to account for some of the basic effects in experimental extinction. As we have seen, it was only partially successful in this context; something like the notion of reactive inhibition seems to have survived experimental test—there does seem to be some temporary inhibition to further responding generated by behavior—but the concept of conditioned inhibition did not survive the experimental onslaught in the study of extinction.

Since, as we have just seen, the decremental effects of massed practice seem to result in both a temporary inhibition (hence recovery as it is shown by reminiscence) and a relatively permanent decrement, several investigators have assumed reactive and conditioned inhibition to be operating; they have tried to find techniques to evaluate unequivocally the role of these inhibitions in the distribution of practice phenomenon (Ammons, 1947; Kimble, 1949a). The logic of the technique for assessing the role of these two inhibitory factors is straightforward. Two groups of subjects are compared, one group working under massed and the other under distributed practice. The difference in performance between the two groups at any point represents the sum of reactive and conditioned inhibition. At some point a rest period is introduced (to allow reactive inhibition to disappear); the gain in performance after rest represents the amount of reactive inhibition that has disappeared, and the difference, after the gain, between the massed and distributed group represents conditioned inhibition.

Despite the simplicity of this analysis, it has many pitfalls. First of all, it assumes that all the reactive inhibition dissipates during rest. Secondly, it assumes that practice in the distributed group was so well separated that no inhibition occurred. Thirdly, it assumes that despite a difference in performance the two groups learned equally (it would not do to have a difference in habit strength as well as in inhibition between the groups).

Finally, it does not take direct account of the warm-up that must occur after a rest. All in all, this is a rather forbidding list of assumptions; so we see that the simplicity of the analysis is more apparent than real.

Some investigators have tried to control or analyze the factors going into these assumptions (Wasserman, 1951), but the results have not always been convincing. Therefore an exact test of the application of formal inhibitory constructs to the massed and distributed-practice phenomenon has not yet occurred. Furthermore, as Underwood pointed out in his previously mentioned series of papers on distribution in verbal learning, a whole host of detailed experimental results cannot be accounted for by these notions.

At the most general level, we can say that there appear to be at least two components to the distributed practice effect, a permanent and a temporary one. Furthermore, neither of these seems to be localized in peripheral structures; they are probably central in origin. Beyond this we cannot specify. It is likely, however, that some of the difference between massed and distributed practice in recall, even when the recall is separated in time from original practice, reflects an underlying difference in what has been learned. In an earlier section we saw that for most complicated tasks, doing is important to learning. Therefore if the decrement produced by massed practice is severe enough to prevent individuals from performing adequately, it is likely that they will not learn adequately. Even so general a notion as this does not allow us to understand in detail some of the effects of distribution. For example, Underwood (1953a) points out that in rote verbal learning by serial anticipation, items anticipated frequently during learning are recalled better after massed practice, while items infrequently anticipated during learning are recalled better after distributed practice. This suggests some kind of interaction between distribution and rate of forgetting. If this is so, then the principles governing the permanent effects of the distribution of practice for material that is easily forgotten will be different from those governing the same effects for material less easily forgotten. Again, this interaction may well account for the greater consistency of findings on the effects of distributed practice on psychomotor skills, since these skills are about equally easy to remember.

The earlier theories and to a certain extent the current theories of the distribution of practice are based upon inadequate empirical exploration. This has not been because there has been a lack of experimental investigation but rather because the problems are complicated. The future course

of theoretical analysis of this problem will require a further systematic investigation of the principal variables.

Information and Effect in the Learning of Skills

There is a famous demonstration by Thorndike (1932b) of the important fact that the consequences, or aftereffects, of an act are important determiners of the future course of that act. If a blindfolded person is requested to draw, with one quick move of a pencil, a line 3 inches long, his accuracy will be poor. If he continues to "practice" without being able to see what he is doing, his performance will not materially increase in accuracy, though it may become less variable. However, if an experimenter measures the drawn line and says "right" if the line is within $\frac{1}{8}$ inch of 3 inches and "wrong" if the line deviates by a greater amount, the subject will rapidly show improvement in his ability to approximate the 3-inch criterion.

This demonstration has been the subject of an enormous amount of misinterpretation. It is meant to show the importance of the aftereffects of a response in determining its direction of change. By implication it is frequently taken to mean that the reward ("right") satisfies the subject's motivation to increase his accuracy and that this has something to do with the improvement. To argue whether or not this is so easily leads us into a morass of slippery concepts. It may well be true (depending upon how systematically defined our notions of satisfaction and motive are), but there is no critical reason why it is necessary to describe this demonstration in these terms. The line-drawing demonstration shows only the operation of the empirical principle of reinforcement, which says that there is a class of stimuli to selectively strengthen every response.

"Right" as a reinforcement is trans-situational (see Chapter 2); that is to say, it will strengthen responses of a wide variety of classes. Almost any sort of information that leads the human learner to make what the experimenter or teacher defines as a correct response is a reinforcement, though some such informational reinforcements are clearly not trans-situational.

Knowledge of results in a training program. A good example of the difficulties involved in getting the correct information to the learner is found in flexible-gunnery training. In firing guns it is sometimes difficult for the trainee to know if and when he hits the target. During the Second World War various training devices were developed to give individuals information about their accuracy in firing. One of these, the Waller trainer, was a screen

which simulated the sky and against which images of planes were thrown; a simulated gun and sights enabled the trainees to fire at the images. The apparatus was rigged so that it would score a "hit" whenever the individual had the sights lined upon the image of the plane and was pulling the trigger. In addition, there was a tone—a "beep"—which was sounded whenever the individual made a hit. This could be thrown in or out of the circuit so that the effect of this piece of information could be studied. Figure 41 tells the story. It is obvious that when the tone was sounded for a hit, the subjects did much better. When the tone was sounded early in practice it resulted in a greater number of hits. Later, when the tone was withdrawn from

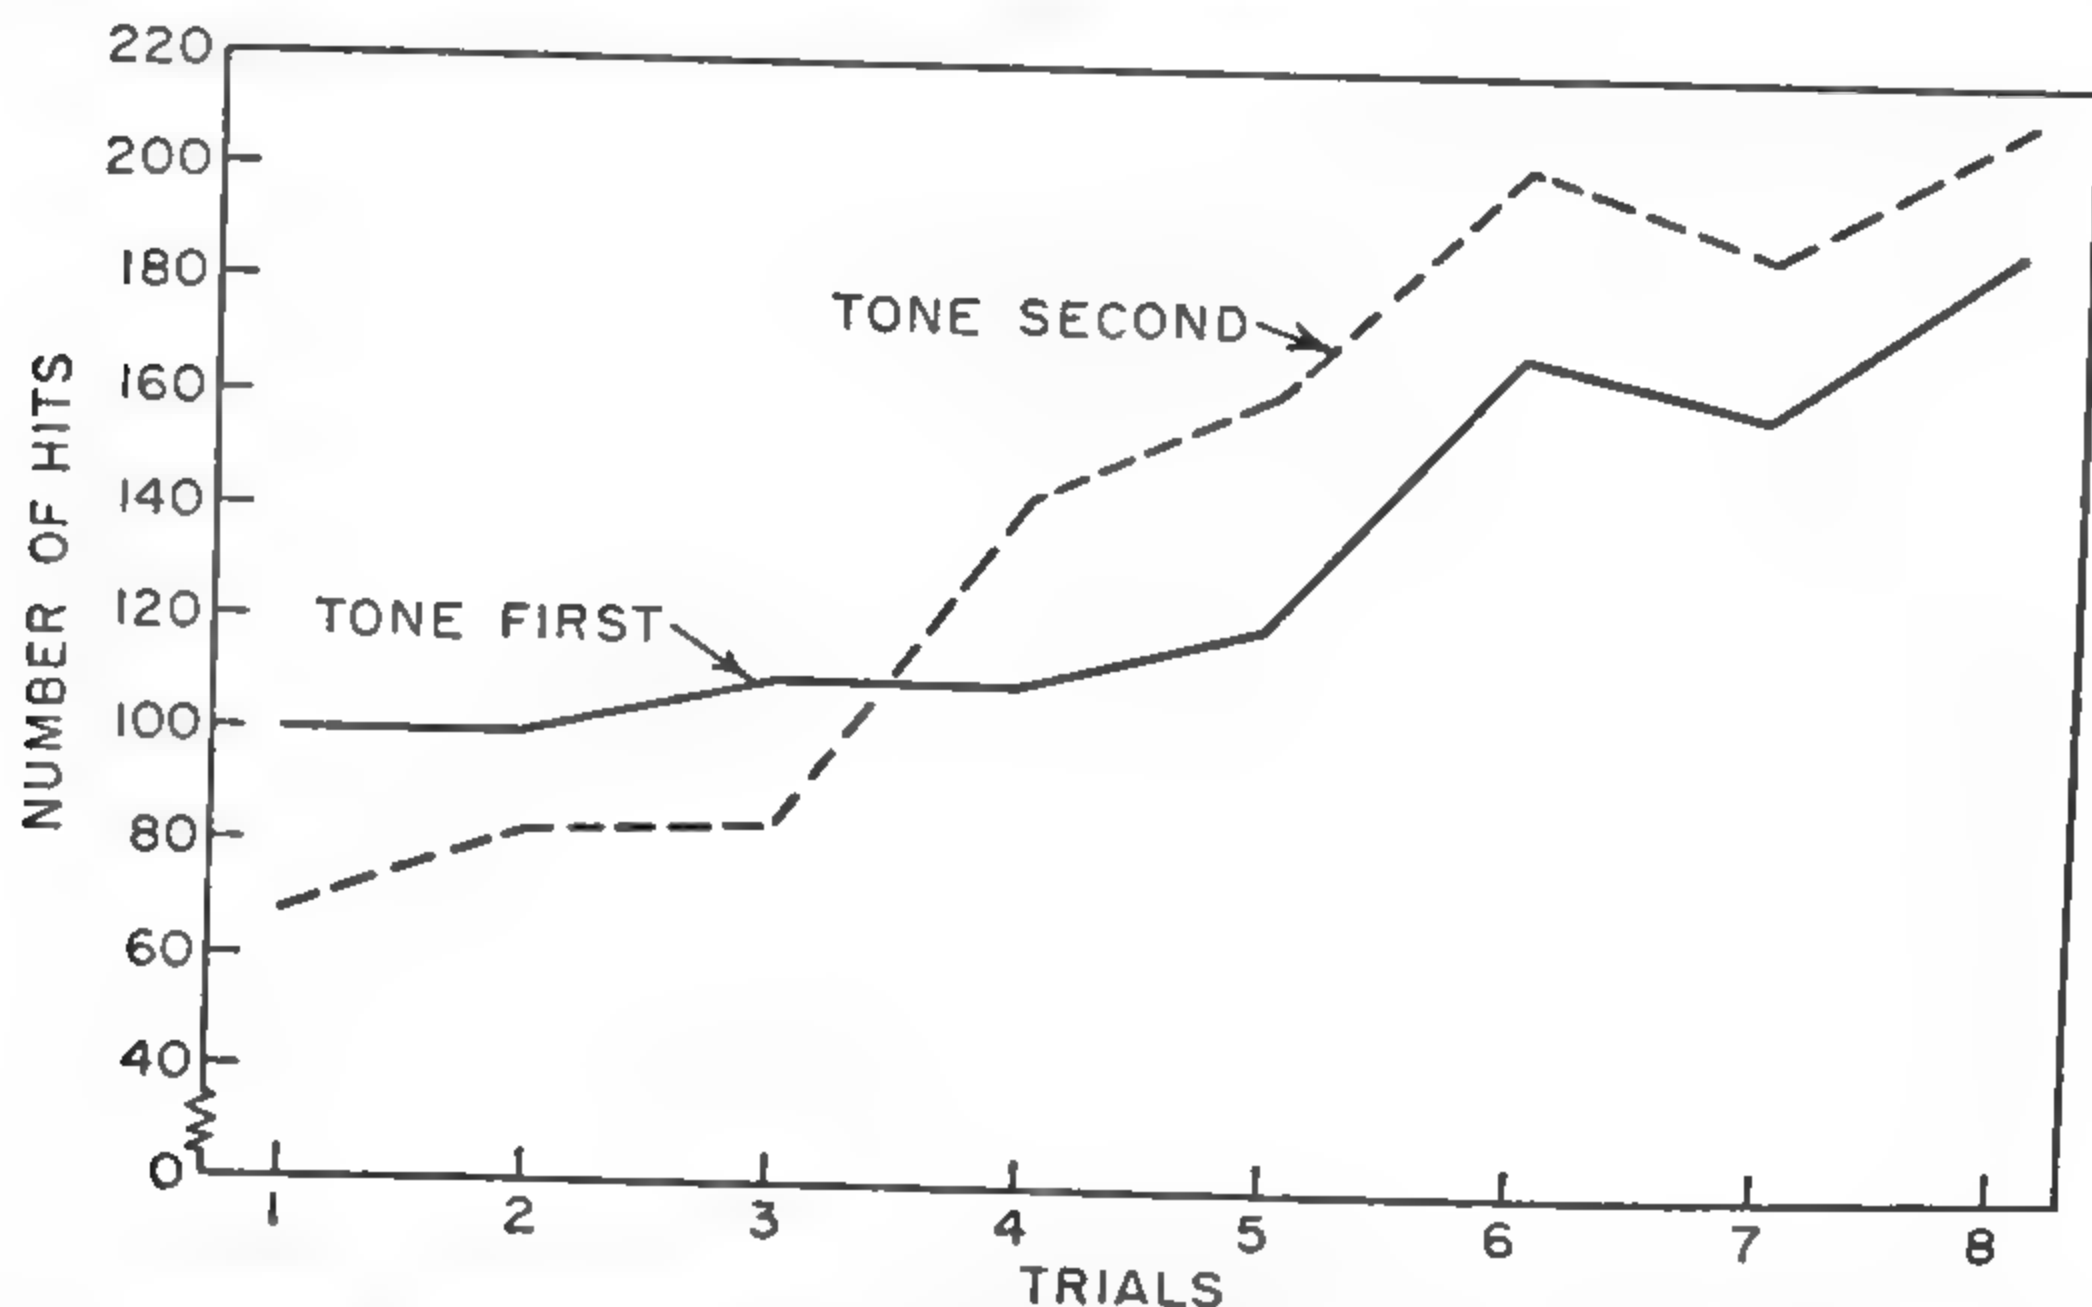


Figure 41. The effect of knowledge of results upon learning to hit a target. Whenever a hit was made, the tone was sounded. One group received the tone at first and not later. The other group received no tone at first but then heard one on the later trials. (Hobbs, 1947.)

this group and given to the second group, the second group immediately increased in number of hits, so that it now surpassed the first group. Unfortunately, the investigators did not include a third group that was never given the tone.

The same group of studies (Hobbs, 1947) shows that where there was little opportunity for knowledge of results, as in the gun camera mounted in the tail of a bomber, relatively little improvement in score took place. This fact ruled out several otherwise excellent, realistic training devices. A second important point is made by the same studies: Sometimes the subject gets the wrong information. Knowledge of results of training with a stationary gunsight and target might be misleading, because in actual practice the gunner must fire at a moving target from a moving plane and must

learn to "lead." Thus knowledge of results on a training device could actually lead to negative transfer on the job. While this may not happen often, it is something which needs to be watched very carefully in any training program.

Delay of information and learning. In examining the influence of reinforcement upon animal behavior we saw that a delay in reinforcement may disrupt behavior and slow down the rate of learning. Is the same thing true of reward and information in adult human learning? At first thought we should probably say no, because the symbolic processes in man would enable him to bridge anything but an enormous delay between the occurrence of a response and the reward or information that comes as its consequence. Let us, however, examine some of the experimental evidence.

In one experiment (Lorge and Thorndike, 1935) subjects learned to toss a ball at a target they could not see. Some subjects were told how well they did on a particular throw immediately after the throw; for other subjects the information was delayed for various short intervals of time. In general, as one might expect, there was little or no difference due to varying intervals of delay in the learning of this task. If, however, the interval between throwing the ball and the information about that throw was filled with another throw, the gain in accuracy was impaired. This result suggests that it is more than simply the passage of time that is important in the effects of delay upon information and reward (we concluded much the same thing in the case of delay of reinforcement on animal behavior).

Saltzman (1951) has demonstrated that a brief delay in reward and information in rote verbal learning can have considerable effects upon the number of errors made by subjects during learning. A delay as short as six seconds increased the number of errors by 50 per cent. Saltzman points out that his subjects rehearsed the response between making the response and the delay of information, and it is quite possible that rehearsing in advance of knowledge of the adequacy of the response made may lead to slower learning. A series of systematic experiments on the influence of such factors is needed.

"Spread of effect" and information. It is quite clear that information about the adequacy of his response allows a learner to adjust his performance more nearly to approximate a norm. Thorndike argued for many years, however, that this was not the only way to look at reward and information. He believed that rewards had automatic effects on responses that were over and above any deliberate use of information on the learner's part; rewards "stamped in" preceding responses or stimulus-response connections.

A series of experiments on a phenomenon called the *spread of effect* convinced Thorndike that rewards acted in this automatic, blind fashion. Thorndike's (1933) original experiments seemed to show that the consequences of a response (reward or punishment) were not specific to the response which they followed but spread more or less automatically to responses temporally adjacent to that response.

The experimental design which Thorndike used was rather unusual. Subjects are given a series of items to which a number of responses are possible, as in a multiple-choice test. The subjects are instructed to choose a response from those available on each trial. Since the subject does not know which responses are supposed to be associated with which word, his responses at the beginning should be fairly random. In one experiment, for example, the subject is read a long list of words. He responds to each word by guessing a number from 1 to 10. The experimenter follows the subject's number with "right," "wrong," or nothing at all. The subject believes that he must learn to associate the right number with the right word. The experimenter, however, has decided beforehand which of the words he will say "correct" to, no matter what the subject gives. The subject is told that more than one number might be correct for any word, so that there is no way in which he can know that there is nothing for him really to learn.

The effect Thorndike found was that the statement "correct" not only increased the probability of repeating the response that it followed, but it increased the probability of repeating the responses before and after the response rewarded. This you can see in a typical spread-of-effect gradient in Figure 42. The responses on either side of the middle response had nothing following them; however, they did increase in per cent repetition over chance expectancy.

At the time Thorndike's results were published, they did indeed seem to give strong support to the notion that the consequences of a response operated by the principle of effect as well as by information. In the intervening years, however, a number of complications have been found in this kind of experiment. First of all, it appears that it is not a stimulus-response connection that is strengthened; it is just the habit of guessing one particular number after some other number (Zirkle, 1946). Even when the order of items given to the subject is changed on every trial, the spread of effect remains.

A number of other experiments combine to show almost overwhelmingly that the spread of effect in this type of experiment is primarily due to the fact that human beings seldom respond randomly, and the effect does not

seem to be related to the principle of effect at all. These experiments show that the guessing habits of subjects are the important determiners of the effect. When no opportunity for guessing habits is permitted, no spread of effect appears (Jenkins and Sheffield, 1946; Jenkins and Postman, 1948). On the other hand, Sheffield (1949) has shown that pure guessing without any information or reinforcement will result in an aftergradient like that found in the spread of effect. So it seems that the dependent probabilities in human verbal behavior provide the background for the spread of effect (Smith, 1949). Finally, Sheffield and Jenkins (1952) have shown that the more closely an experiment approximates Thorndike's original one without

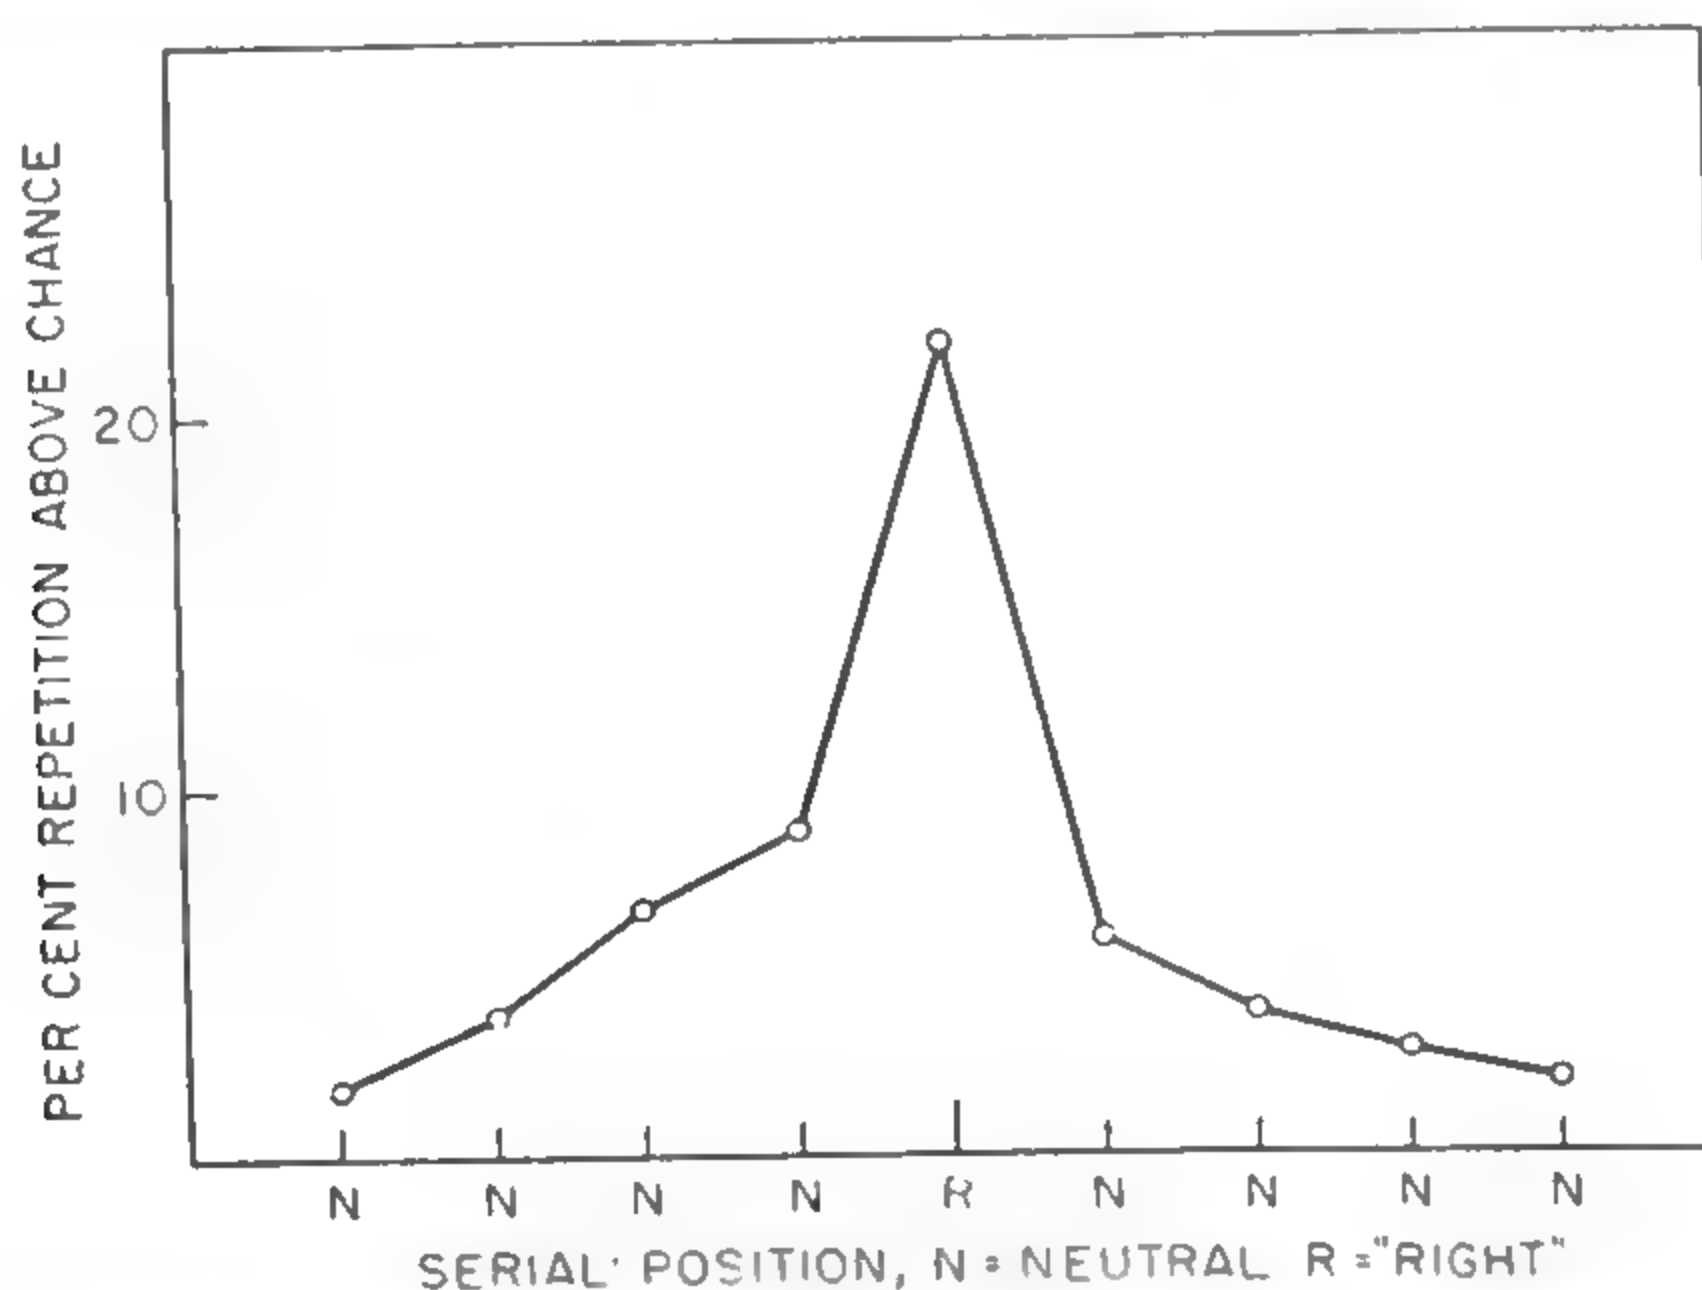


Figure 42. The spread-of-effect phenomenon. The probability of items adjacent to a rewarded item in a series is increased, as illustrated in this hypothetical example.

the presence of reward, the closer one gets to a typical spread-of-effect gradient. Thus there seems to be little doubt that Thorndike's original spread of effect was an artifact.

Some investigators have pointed out the possibility of another, more general kind of gradient (Marx and Bunch, 1951; Marx and Bernstein, 1955). This is not based upon temporal proximity to the reward, but upon resemblance of other items to the rewarded item. Such a gradient would appear to be reasonable under certain conditions, but it is not the same as spread of effect, and it adds nothing to the notion of the automatic strengthening effect of reward. Rather a reward appears to function (1) to keep activity going (as it does for animal behavior) and (2) to provide information about the adequacy of past behavior and the direction future change of behavior should take.

Other Conditions of Practice

Of course there are a large number of factors that determine the efficiency of learning for particular tasks in formal training programs or in schooling. Sometimes these factors are so highly specific to particular tasks that they are of little general interest in the psychology of learning, and sometimes they are unsystematic arrangements of more basic variables. For example, consider the problem of active participation in learning.

It is no surprise to anyone that active participation in a task yields superior performance, recall, or transfer for that task. For example, interruption of an educational film for multiple-choice questions increases the effectiveness of the film as a teaching device (Kurtz, Walter, and Brenner, 1950), though perhaps, rather surprisingly, no more so than simply repeating the film a second time. Obviously, the questions provide more active participation on the part of the viewers than does merely watching the film.

When we ask why active participation is more effective than passive participation, we immediately see that it is a complicated question. Participation varies for different tasks, and it must have somewhat different effects. It may supply the learner with knowledge of results with short delay, increase the learner's motivation, or help him eliminate errors early in practice.

What is surprising, perhaps, is that retention of verbal material is better when time is taken away from reading or study of the material and put into active recitation. In one study (Gates, 1917) subjects who spent up to four-fifths of their total study time in recitation rather than in reading showed the greatest improvement. The advantage gained from recitation is frequently large, and this is one of the points about habits of study that is emphasized in manuals on study habits (see Robinson, 1946). Very likely recitation is important to good study habits because it provides direction to the learner's work and readily supplies information or knowledge of results about performance.

Another problem that sometimes arises in application of principles of learning to formal training or schooling is the question of whole and part learning. This originally arose in the context of rote memorization, where perhaps it makes most sense, and it was concerned with whether it is more efficient to memorize by reading the material from beginning to end each time or by breaking the material into its constituent parts and studying each separately.

A little reflection will show that even in the limited context of rote

memorization, the question of whole and part learning is complicated. What, for example, is the role of length of task? Perhaps for short tasks, such as memorizing a sonnet, the whole method would be better, while for longer tasks, memorizing a book from *Paradise Lost*, for example, the part method might be more efficient. Unfortunately we do not have systematic information to answer the question about the relative efficiency of whole and part learning for all tasks. We cannot even make a reasonably sure guess, as we might in the case of massed versus distributed practice. We do know that some conditions favor whole learning and others favor part (Jensen and Lemoire, 1937), but we do not know enough about the interaction of this problem with other methods of practice to make clear recommendations about a wide variety of tasks.

TASK VARIABLES

Let us now consider how some of the important differences between various tasks affect learning. No task that a human being learns is quite like any other. Some of the differences between tasks are trivial and others are of limited general interest. There are, however, a number of ways in which tasks differ that are fundamental to the psychology of learning.

First of all there is the important difference between those tasks that are primarily verbal and those that are not. Then there is a dimension of perceptual emphasis along which tasks vary; some are primarily a matter of learning to make discriminations and to form new perceptual organizations. Other tasks are primarily a matter of adjusting motor performance to meet some norm of skill. Still others are matters of conceptual organization. In tasks that demand simple memory, there is an important dimension that runs from demanding memorization of a collection of unstructured elements to demanding memorization of highly organized and meaningful material. Compare, for example, an attempt to memorize a list of 100 nonsense syllables with an attempt to memorize 100 successive words from a story in a popular magazine.

Then there are the two important dimensions which run through practically all types of tasks—difficulty and length, or amount, of material to be learned. Though we regard these as fundamental dimensions, they are not necessarily pure. Task difficulty is frequently correlated with some other way in which tasks vary (meaningfulness, for example), and it is by no means a simple variable. As a matter of fact, such variables as task difficulty are probably only rough empirical classes, and the future development

of learning theory may enable us to dispense with such notions altogether. For now, however, a considerable part of our understanding of the basic mechanisms in the psychology of learning is tied up with such variables as difficulty and amount, and we must regard them carefully.

Difficulty

A difficult task is generally one that takes us a relatively long time to learn. Yet, it is not possible to equate difficulty with the length of time or the number of trials it takes for a learner to reach some arbitrary criterion, since other conditions influence the length of time spent in learning. The abilities of the learner, his motivation, the amount of preceding practice, and the distribution of work are all factors that will influence the length of time it takes to reach some criterion of performance. We can, however, define a change in task difficulty when we keep all conditions, such as distribution of work, motivation, ability, constant and vary some aspect of the task in such a way that the time taken to reach a criterion is changed.

Difficulty and Response Characteristics

Number of available responses in trial-and-error learning. Trial-and-error learning generally refers to those situations in which learners cannot make use of a general principle or of previous knowledge, and in which they must try to find the correct solution by exploring the available alternatives. In such learning an important determiner of difficulty is the number of choices the learner is given (Noble, 1956). The larger the number of alternatives (assuming the number of correct alternatives remains the same), the greater the difficulty of the task. For example, in one experiment (Noble, 1955) subjects had to learn to push four buttons in a particular sequence; in this situation the number of trials required to reach an arbitrary criterion was enormously increased when the number of available buttons increased. If only four buttons were present, they were all relevant to the task, but when the buttons were increased to 10 in number, 6 of them were irrelevant and the subjects had to learn to eliminate these.

Interaction between number of available responses and task length. The influence of number of alternatives on performance is modified by the length of the task. This is demonstrated in a series of studies by Brogden and Schmidt (1954a, 1954b). These investigators studied the effects of increasing both the number of response alternatives available at each choice point and the total number of choice points. They used a verbal maze, one in which the learner must find his way by guessing which of a number of

available responses at each choice point will allow him to go on to the next choice point and eventually the end of the maze. In their experiments, Brogden and Schmidt studied two different lengths of maze, one 16 units long and another 24 units long.

The results of this study can be seen in Figure 43. Both the number of errors and the time taken to reach a criterion score increased linearly with an increase in the number of alternatives. It is interesting to note, however, that for the two mazes of different length, the rate of increase in time and errors with increasing number of alternatives is different. Apparently, the longer the maze the greater the difficulty added by increasing the number of alternatives at each choice. A careful examination of the curves in Figure 43 will reveal that the increase is disproportionate in the case of the

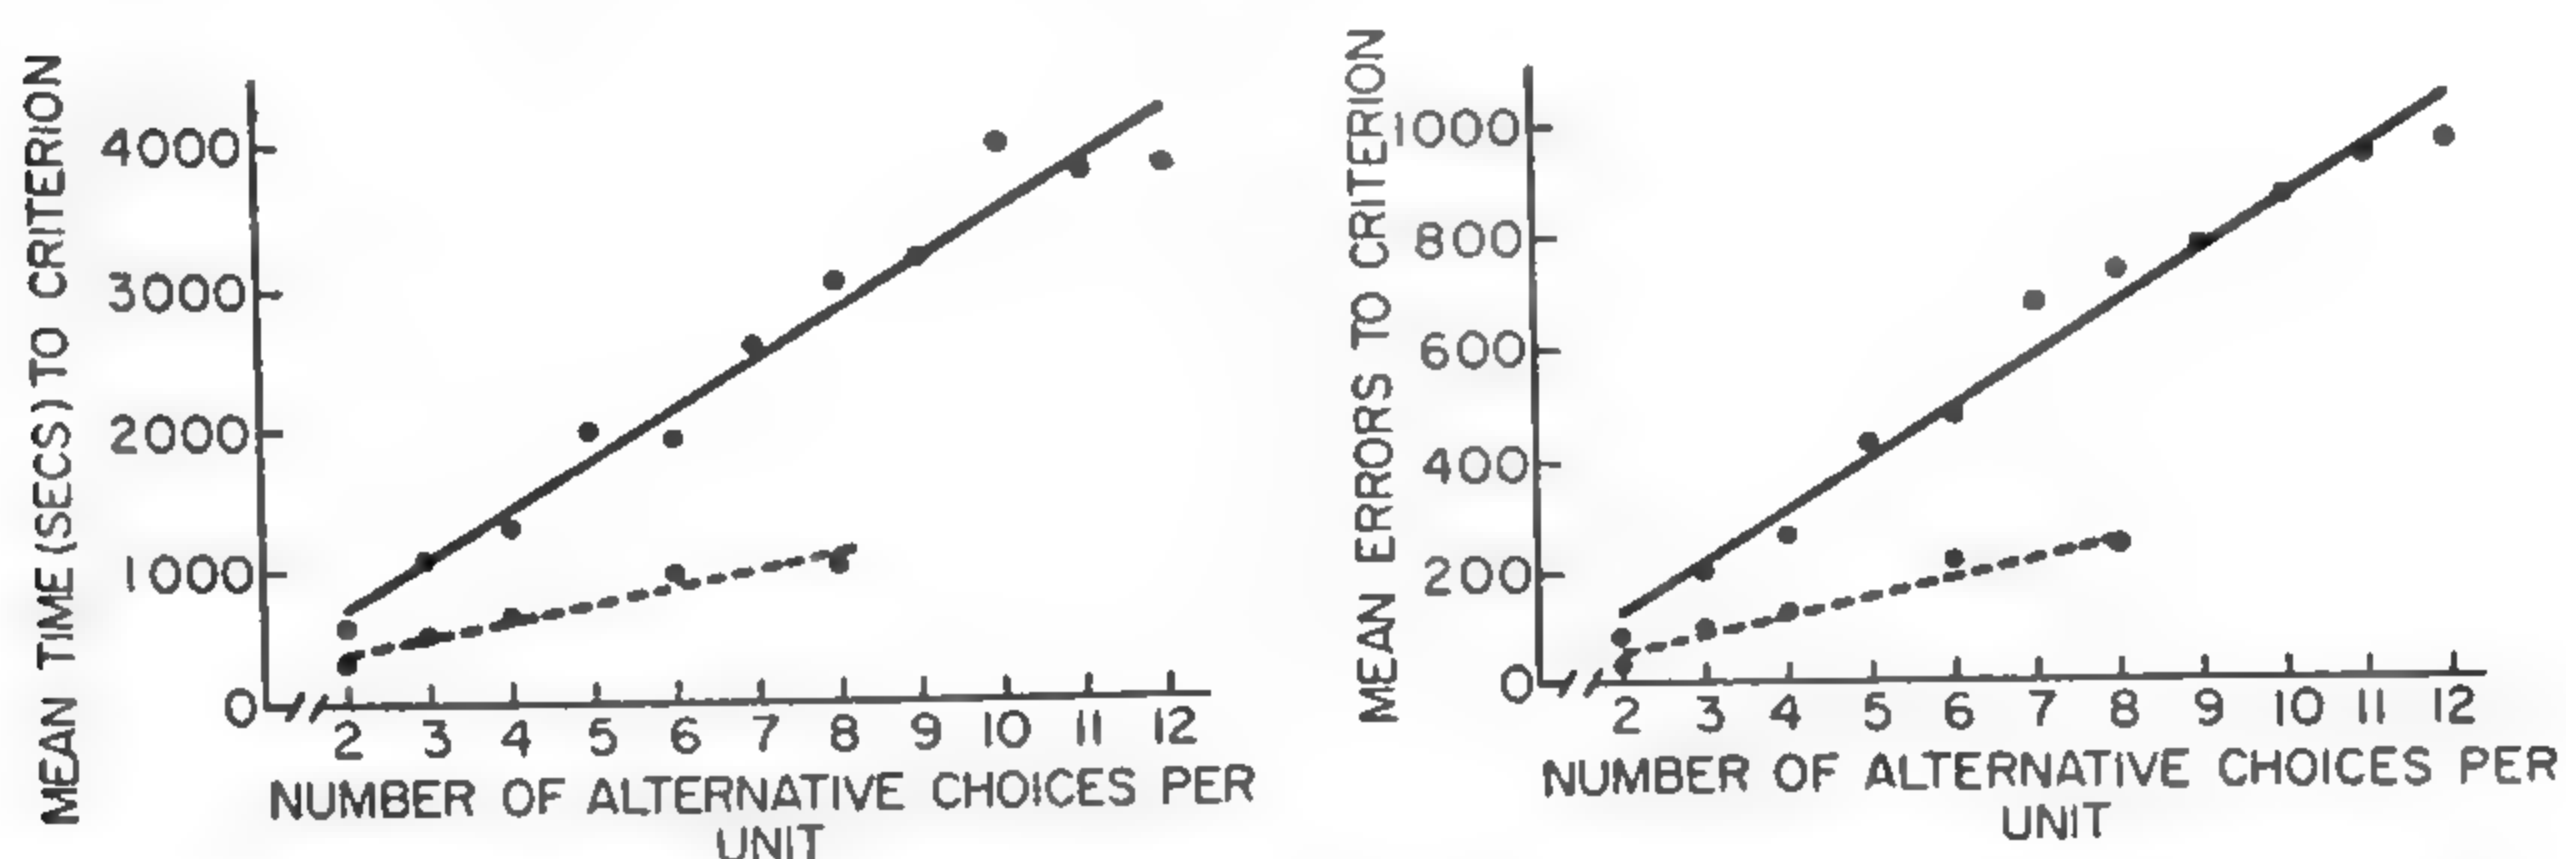


Figure 43. Curves showing the increase in errors and time to criterion with an increase in the number of choices at each choice point in verbal mazes. The dotted curves are for mazes 16 units long and the solid ones for mazes 24 units long. (Brogden and Schmidt, 1954b.)

longer maze. In other words, the greater time and greater number of errors to criterion is not simply accounted for by the fact that there are more choices to be made; increasing the number of choices and increasing the length of the maze seem to have different effects.

In these studies one important feature turned up that is of critical interest. In general, increasing the number of alternatives at each choice did not increase the number of trials required to criterion. Thus, while the total number of errors made by subjects and the total time to criterion were increased by adding to the number of choices, the subjects did not increase the number of times they had to go through the maze in order to reach criterion. This is in contrast with the results found by Noble (1955), and the answer lies in the fact that Noble's definition of trials differs from that of Brogden and Schmidt. In Noble's study trials means the number of guesses the subject makes; Brogden and Schmidt define trials as the number of times the sub-

ject had to go through the different choice points. Miller (1956a) has pointed out that the reason trials to criterion was not influenced in Brogden and Schmidt's study was that subjects could organize the responses available so that they did not need to take into account all the possibilities. Once the correct choice is found, memory for this choice is not influenced by whether there are 4 choices or 10. What does make a difference, as we shall shortly see, is the total number of things it is necessary to remember. The number of alternatives will influence the total number of errors subjects make (and therefore the total time, since errors take time), because the greater the number of alternatives the lower the probability of making a correct guess. If I ask someone to guess in which hand I have a penny, he has a better chance of guessing correctly than if I had asked which of eight pockets on my person contains a penny. Since, in Brogden and Schmidt's study, the subject must eventually guess the correct alternative before going on to the next choice point, the total number of times the subject goes through the sequence of choice points is unaffected.

Therefore, while increasing the number of alternatives will increase the number of errors made during learning, it will not increase the difficulty of remembering the correct response.

Difficulty and Characteristics of the Material Learned

The relationship of characteristics of a task to be learned and difficulty of learning is a fairly broad question. We can ask, for example, What is more difficult to learn, typing or driving an automobile? Such a question would have little general value, since these two tasks differ from one another in a large number of characteristics. If, however, we compare tasks which differ in only one characteristic we can learn something about how variation contributes to the difficulty of tasks.

Association value. One important characteristic of verbal material is its richness of contextual association. Some words have a greater potential for eliciting associations than other words. Even nonsense words or syllables differ enormously in their potential for eliciting associations. Various lists of nonsense syllables have been carefully examined for their association value, and it is possible to choose syllables of different value for different kinds of experimental studies.

Noble (1952a) has made use of an index of meaning applied to two-syllable nonsense and ordinary English words. This index was obtained by counting the average number of associations given during a 60-second

period to each word. The obtained index ranged from 0.99 for the word GOJEY to 9.61 for the word KITCHEN. It is interesting to note that such a nonsense word as ROMPIN has a higher derived index of meaning than such infrequent but real English words as ICON, MATRIX, and BODKIN. Noble (1952b) used these words in a study of rote learning by serial anticipation. It turns out that a list constructed of words with an average meaning index of 1.28 took almost three times as many trials, on the average, to learn as a list constructed of words with an average meaning index of 7.85. These results, taken together with earlier studies of three-letter nonsense syllables (McGeoch, 1930), make it clear that associative potential of individual verbal items is an important determiner of ease of learning. It is worth noting that associative potential plays essentially the same role in paired-associate learning (Noble and McNeely, 1957).

Similarity. A variable of considerable importance from a theoretical point of view is that of similarity among items in the material to be learned. Similarity, as we shall see in later chapters, is a rather complicated dimension. Any task consists of stimulus and response components. Both may vary in similarity among the separate items and both may vary with respect to each other. In rote serial learning the same item serves as both stimulus and response in succession, while in paired-associate learning there is a clearer separation between stimulus and response elements. Therefore, we might expect some differences to exist between the effects of variation in stimulus similarity and in response similarity. Hence, tasks in which there are different relationships between the items must be considered separately.

If we consider the rote serial learning of verbal items we see that an increase in the similarity between items generally results in an increase in the number of trials to criterion. For example, Underwood and Goad (1951) showed that lists composed of adjectives highly similar to one another, such as "elated," "gleeful," "carefree," took more trials to learn than lists composed of unrelated adjectives like "worldly," "fiery," "blotchy." The same thing is apparently true of nonsense syllables, in which the similarity is achieved, not by meaning, but by duplication of letters (Lazarus, Deese, and Hamilton, 1954; Underwood and Richardson, 1956).

Underwood (1951) has also studied the influence of similarity upon paired-associate learning. In his study he compared learning of lists composed of distinctly different adjectives with learning of lists in which similar adjectives were used as stimuli and responses. In general, the more similar the adjectives the larger the number of trials to criterion. Unfortunately,

Underwood did not conduct studies in which he varied stimulus similarity while holding response similarity constant or varied response similarity while holding stimulus similarity constant. From other studies, some of which we shall examine later, we can conclude that probably both stimulus similarity and response similarity increase the difficulty of a paired-associate task. High stimulus similarity produces high stimulus generalization from one item to another; since the stimulus items are less distinctive, they should be harder to discriminate. Likewise, if response items are highly similar, learners should have a harder time in finding and remembering their distinctive properties.

In general, it is probably safe to conclude that *within* a given task, high similarity of stimulus and/or response components increases the difficulty of that task. We should note, however, that this is only necessarily true where the learner is required to discriminate among all the components of the task. If, for example, any one of several responses will do in response to a given stimulus, response similarity could well help rather than retard learning. This case illustrates the essence of the difference between negative and positive transfer, and we shall have many opportunities to examine such cases in the next few chapters.

Amount to Be Learned

We hardly need an experimental study to tell us that the more material a subject has to learn, the longer the time required, or indeed, that a greater amount to be learned requires more trials for learning. The exact nature of the relationships between time and trials necessary for learning and amount of material to be learned is interesting. It is also somewhat surprising, and it will repay us to examine it carefully. A classical monograph by Lyon (1917) summarizes much of the best available evidence on these relationships, and we shall draw heavily on Lyon.

Amount to Be Learned and Time per Item

Nonsense syllables. Perhaps the most meaningful way to present the relationship between amount to be learned and time to learn is not to compare the total learning time with the number of items for different numbers of items, but to see how length of time to learn increases as items are added. Such a comparison can be seen in Figure 44. This figure shows the results of two of Lyon's experiments on the time required to learn lists with different numbers of nonsense syllables. In one experiment Lyon learned nonsense syllables by continuous practice, and in the other by repeating

the list only once each day. The curves show the *time per syllable* as a function of number of different syllables. As we can easily see, the time per syllable increases rapidly with continuous practice. We can imagine, from our earlier discussion, that this very rapid increase in time per syllable probably represents something more than the inherent difficulty of learning longer lists; this may represent the cumulation of response-produced inhibition as well. So, for our present purposes the curve for one trial per day is probably more meaningful. It is rather surprising that while there is an increase in the time per syllable with longer lists, it is not a very

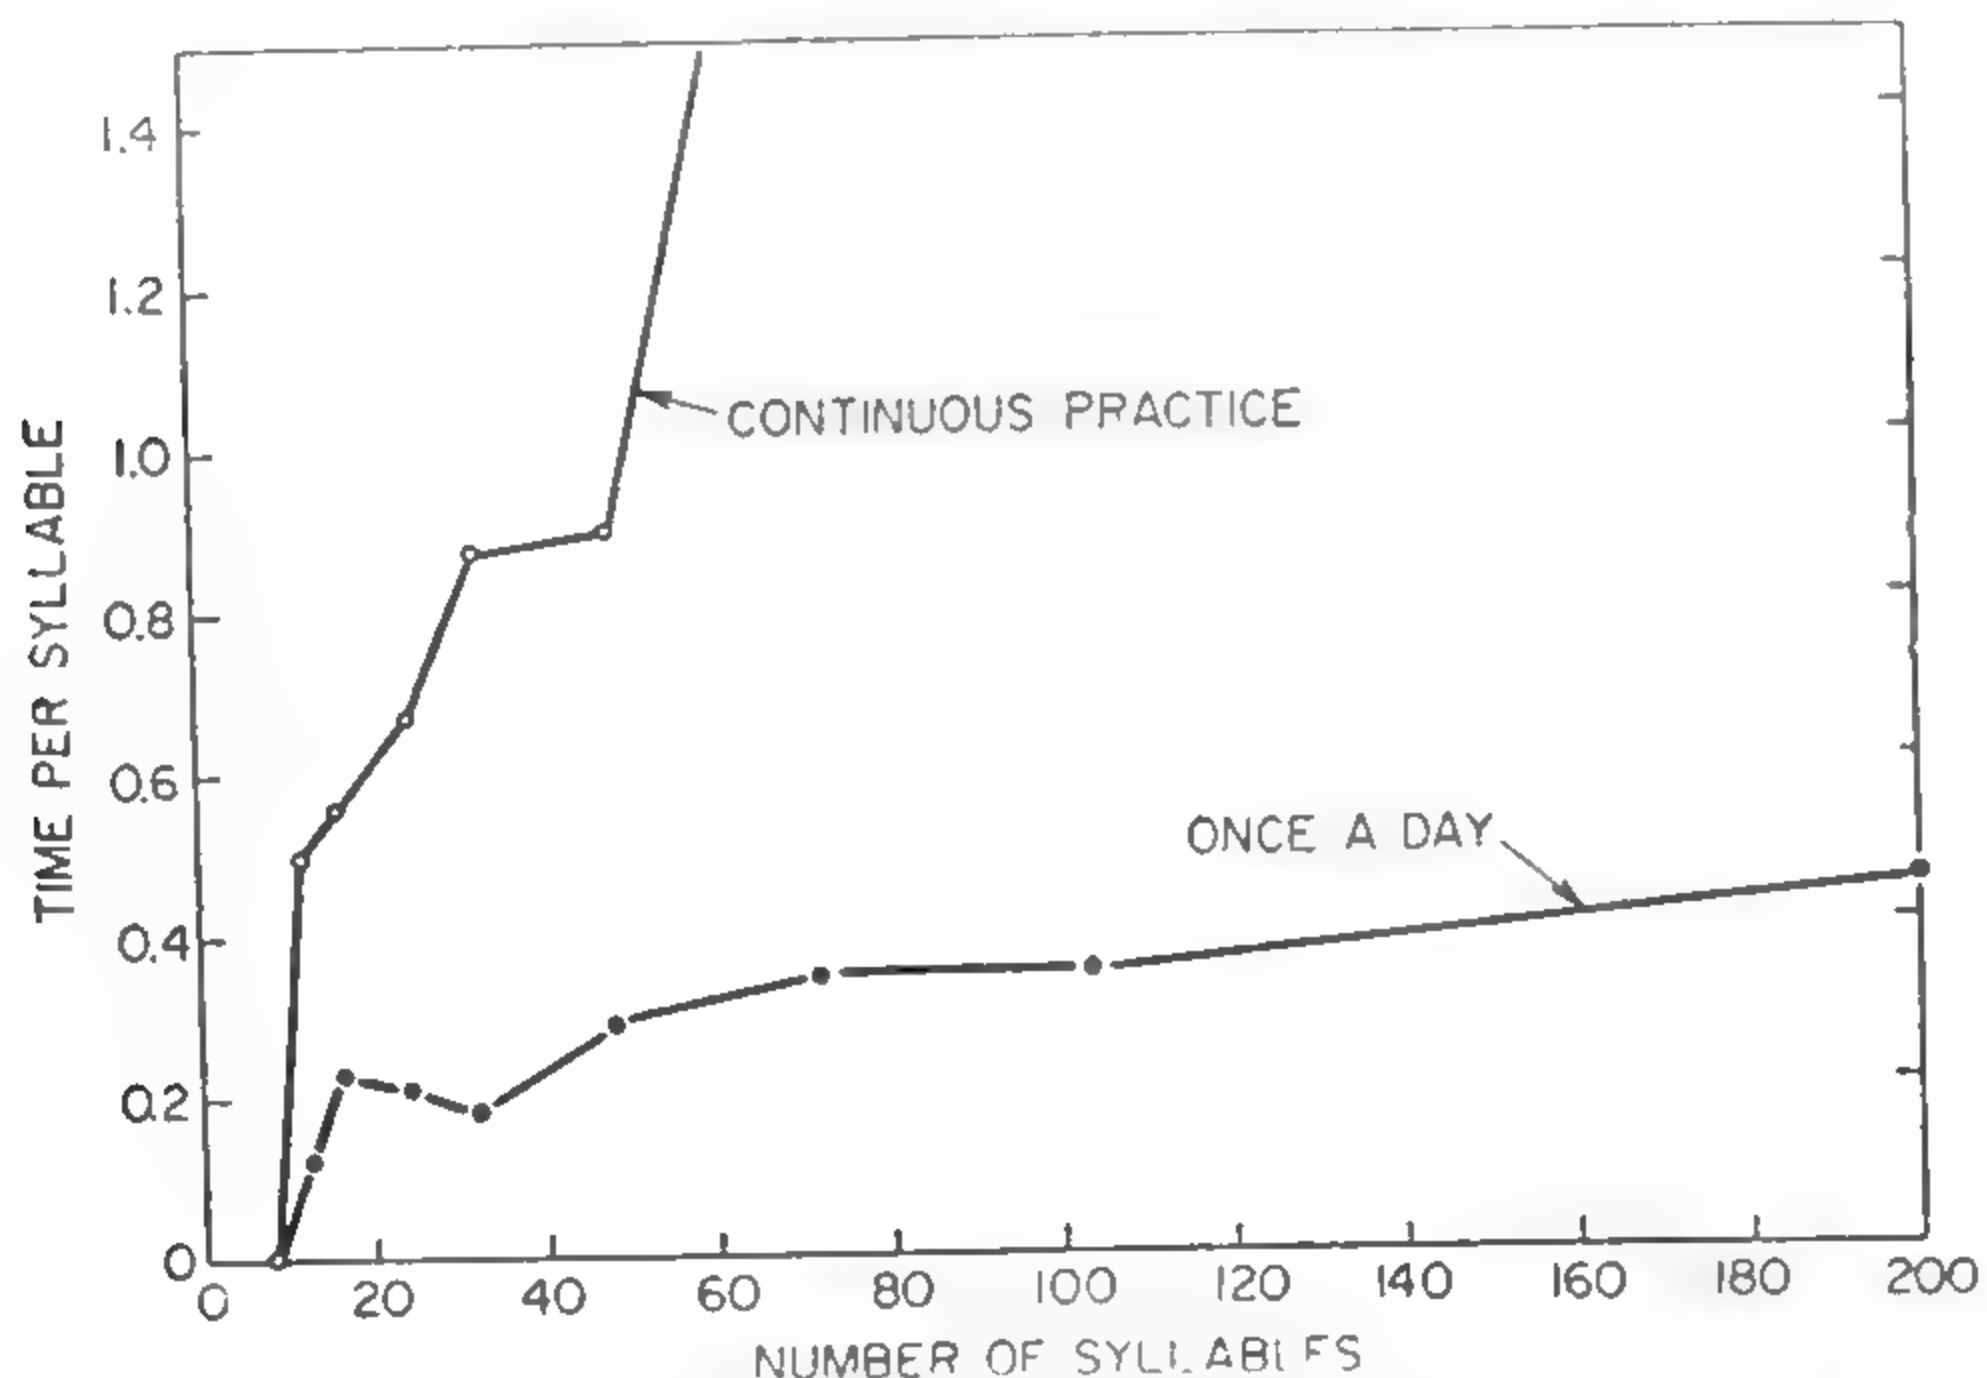


Figure 44. The time per syllable required to learn lists of nonsense syllables of varying length for distributed and massed practice. The data for these curves were obtained on one highly trained subject. (Data from Lyon, 1917.)

great increase and appears to level off. That is to say, with small numbers of syllables the addition of a few more syllables makes a large difference in time per syllable, but with a large number, the addition of the same number of items makes little difference in time to learn.

Therefore, if a large number of items are to be learned by the method of serial anticipation with distributed practice, the addition of an item or two costs little more than the additional time required to read the syllable. This result is somewhat surprising, because intuitively one might expect that if one has a great number of syllables to keep in mind already, the addition of one more would indeed add to learning time. Thurstone (1930b) has shown that it can be derived from his general learning equation that time per item should increase as the square root of the number of syllables beyond

the immediate memory span.² Such a conclusion is in accord with the data shown in Figure 44 for the once per day method as well as data from other, earlier investigators.

Meaningful material. Lyon also studied the amount of time that it took him to memorize poetry. The results with poetry were much like those he found with nonsense syllables. While additional material did require an additional amount of time per item, each item took less and less additional time. The items in this case were stanzas of poetry. Of course, the material is much longer, but each syllable of the poetry took less time to memorize than each individual nonsense syllable.

If the meaningful material were completely redundant, of course each item would require no additional time to learn. Suppose, for example, I said that I would like someone to memorize a list of digits. If the digits were arranged as follows, 5-8-4-6-5-8-2-1-5-4-9-3-1-4-9, etc., he would have to memorize them by rote, and the more items the more time it would take. If, however, the digits were arranged as follows, 3-2-1-4-3-2-5-4-3-6-5-4-7-6-5-8-7-6-9, etc., he would easily get the principle and additional items (assuming they followed the principle) would take no increment in time; the time per item would be the same no matter what the length of the list. Such a list would be a redundant list; one could reconstruct the entire list from the simple sequence 3-2-1-4-3-2. Now all meaningful material is to some extent redundant; therefore the additional time per comparable item should be less for meaningful material than for some nonredundant material like random lists of nonsense syllables.

Amount to Be Learned and Number of Repetitions

Nonsense syllables. If the exposure time of each item in a verbal learning experiment is carefully controlled, the number of repetitions would be exactly proportional to the time. In Lyon's experiments, however, he did not control exposure time, so that the correlation between time of learning and number of items to be learned is not perfect. Indeed, when we look at the results of Lyon's observations on the amount to be learned and number of repetitions shown in Figure 45 we see that number of repetitions and time per syllable behave very differently. The curve in Figure 45 shows the number of repetitions, or trials, per syllable necessary for learning from the same study as the distributed practice example in Figure 44. The number of repetitions required per syllable increases rapidly at first, but then,

² The immediate memory span is the number of items that can be immediately recalled after only one trial.

rather surprisingly, it reaches a maximum and begins to decline. The result is that Lyon required *fewer* repetitions per syllable to learn a list of 200 nonsense syllables than he required to learn a list of 16 nonsense syllables (a little multiplication will show, of course, that the *total* number of repetitions was greater in the case of the 200 syllable list). These results are not peculiar to Lyon as a learner; in his monograph he summarizes results of

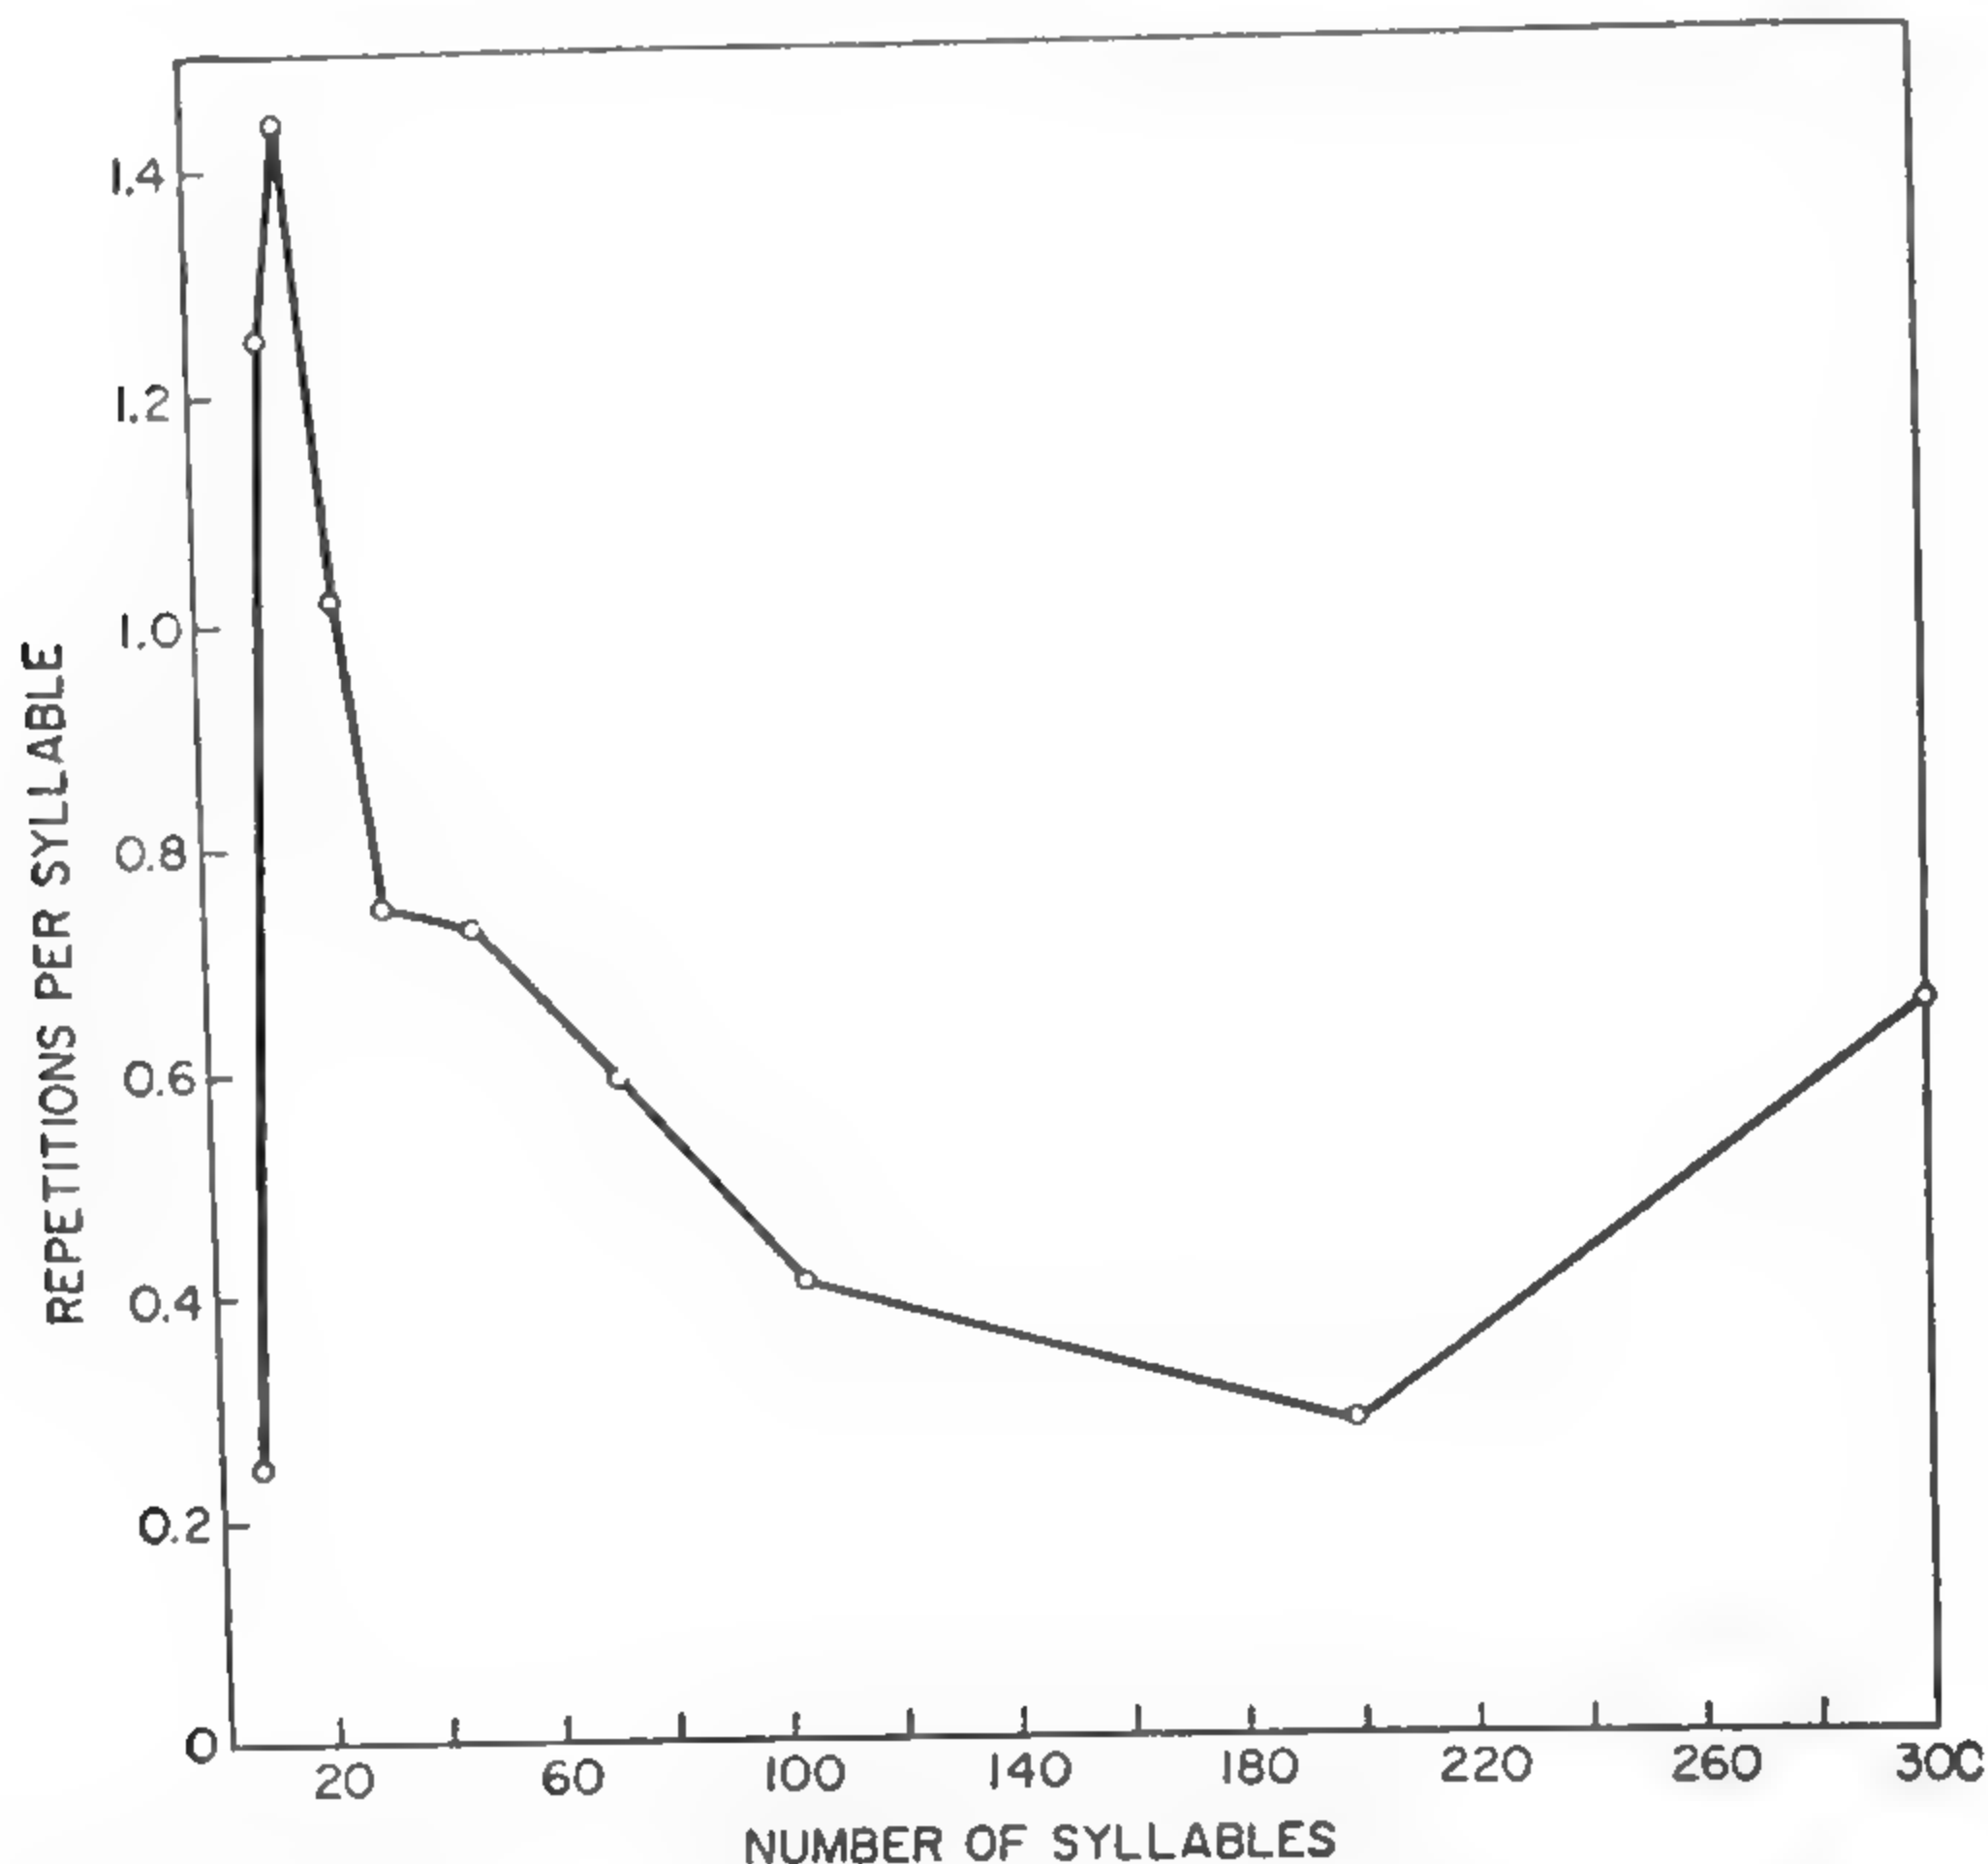


Figure 45. Number of repetitions per syllable required to learn lists of nonsense syllables of varying lengths by the once per day method. This curve is based upon the same study illustrated in Figure 44. (Data from Lyon, 1917.)

experiments by Ebbinghaus and Meumann which show essentially the same thing, though not in so dramatic a form as Figure 45.

Thus it seems probable that the number of repetitions does not increase with the number of syllables beyond some relatively small number, and that it may actually decrease, especially if the exposure time is uncontrolled as it usually was in these early experiments. What if the exposure period is carefully controlled? Unfortunately, we do not know what happens, since there are no data over a range of material comparable to that investigated by Lyon in which exposure time per item is controlled. These experiments

are difficult and tedious, and no modern investigator has tackled them on the scale that Lyon did. It is quite possible that repetitions per syllable would not decrease if this were done, but a comparison of Figures 44 and 45 suggests that it is unlikely that the repetitions per syllable would actually increase much with the number of syllables. The net result of these studies is to make it probable that the increase in amount of work (time or trials) to criterion is not disproportionate to an increase in amount of material, except possibly in the transition from the immediate memory span to longer lists.

Meaningful material. Cofer (1941) has studied the relationship between amount of material and number of trials to learn with meaningful verbal material. As in the case of Lyon's experiments Cofer did not control the exposure time; consequently we should expect his results to be much like Lyon's. And in fact they are, when the differences produced by Cofer's more unpracticed subjects, etc., are taken into account. There is an absolute increase in the number of trials required to reach criterion as more material is added, but the number of trials per unit word actually decreases with longer lists. This incidentally is true whether one considers number of trials per unit word, per unit sentence, or per unit significant idea, and it is true whether instructions emphasize verbatim (rote) learning or logical (ideational) learning. Of course, the absolute number of trials per word is much smaller than in the case of Lyon's experiments on nonsense syllables, even though Cofer's subjects were relatively unpracticed. This is, of course, because of the higher redundancy in meaningful material.

The results of the available experimental evidence clearly suggest that amount of material does not necessarily disproportionately increase difficulty of learning as has commonly been supposed both by learning theorists and those only casually interested in learning. This conclusion, however, is limited by the word "necessarily" because many of the important experimental parameters have not been thoroughly explored. It is possible, for example, that material with high similarity between items would result not only in an over-all greater number of repetitions per syllable, but repetitions may be disproportionately greater for longer lists composed of such material. This whole question of the relationship between trials to criterion and amount of material is ripe for exploration, particularly since there is a very real possibility that the results of such exploration may lend themselves readily to interpretation in terms of measures of information.

CHAPTER 9

TRANSFER OF TRAINING

There is no more important topic in the whole of the psychology of learning than transfer of training. Nearly everyone knows that transfer of training is basic to educational theory. Practically all educational and training programs are built upon the fundamental premise that human beings have the ability to transfer what they have learned from one situation to another. This assumption is evident in things as different as the curriculum of a college of classical tradition and an adolescent who is persuaded by his mother to attend dancing school in the fond hope that training in the graces of ball-room dancing will permit him to navigate the length of the house without danger to less sturdy furniture.

The problem of transfer of training arose in the context of educational theory, but it has always been fundamental to psychological theory as well. Indeed, the basic psychological problems in the transfer of training pervade the whole of the psychology of human learning, and in one sense the majority of the topics this book examines are problems in transfer. Thus, the stimulus generalization of a conditioned response, the chaining of responses in maze learning, and the retention of verbal learning all involve special cases of transfer of training.

In this chapter we shall examine some major techniques for the study of transfer of training as a problem in itself, some conditions that lead to positive transfer between tasks, and some theoretical principles that arise out of the study of transfer.

THE MEASUREMENT OF TRANSFER

Transfer of training means that training or performance on one task has influenced performance on another task. Thus the fundamental technique

in the study of transfer of training is to compare the performance on a particular task of two groups of subjects, one of which has practiced on another task. There are a number of ways in which such an experiment can be performed and the comparison evaluated. These will be discussed in the following section.

Experimental Designs in the Study of Transfer

The proactive design. The *proactive design* is the simplest arrangement in the study of transfer. Basically, there are two groups of subjects, one of which practices a preliminary task and both of which then are tested on a subsequent task. Thus the basic conditions can be summarized as follows:

Experimental group

Learns task A Test on task B

Control group

Rests Test on task B

If task A has had an effect upon task B, these two groups should perform differently on the test of task B. If the experimental group performs significantly better than the control group on task B, we can say that proactive facilitation, or more generally positive transfer, has taken place. If, on the other hand, the experimental group performs significantly worse on task B, we can say that proactive inhibition, or negative transfer, has taken place. If there is no difference between the two groups on task B, of course, no transfer has taken place.

It takes at least two groups of subjects to perform a transfer experiment of this type. For many purposes, however, we may wish to increase the number of comparisons. We may have, for example, several variations in task A. In such a case we should want to compare each of the variations in task A with a control or rest condition. Furthermore, we may want to compare the amount of transfer between task A and B with the effect of direct preliminary practice on task B. Thus a more complicated proactive experimental design might be as follows:

Group 1: Learns task B Test on task B

Group 2: Learns task A₁ Test on task B

Group 3: Learns task A₂ Test on task B

Group C: Rests Test on task B

It is surprising to find out that there are circumstances when transfer from one task to another is greater than the advantages of direct practice on the task to be tested. We shall examine some of the circumstances surrounding this unusual case later on.

One further variation in the proactive design provides for a retention test for the second task. In the simplest example of this design, one group learns task A, then both practice on task B, and finally both are tested for retention of task B. The design is as follows:

Experimental group

Learns task A Learns task B Tested for retention on task B

Control group

Rests Learns task B Tested for retention on task B

This particular design is not so important to the theory of transfer of training itself. It is, however, fundamental in the application of the theory of transfer to an understanding of the forgetting process, and in the next chapter we shall have ample opportunity to see examples of the use of this design.

The retroactive design. The other commonly employed experimental plan for the study of transfer is known as the retroaction experiment. In this design, an experimental group practices first on task A, then on task B, and finally on task A again. The control group practices only on task A for a time comparable to the experimental group. Thus the experimental design is as follows:

Experimental group

Learns task A Learns task B Test on task A

Control group

Learns task A Rests Test on task A

The critical comparison is on the final test of task A. For readily apparent reasons, this design is sometimes called the "fore-and-after" method. It has the advantage that the initial practice on task A can serve to equate the two groups of subjects in performance and thus reduce extraneous variability on the final test. Furthermore the design enables one to examine directly the influence of learning task B upon the *retention* of task A, and this is sometimes very important.

The control group in proaction and retroaction experiments. In the ideal experiment the control subjects would be held in absolute inactivity, so that the resulting differences between the experimental and control conditions could be entirely attributed to the task assigned the experimental group. In practice, of course, it is impossible to keep human or animal subjects completely inactive. This seldom is a serious matter with animal subjects, however, since we have good grounds for supposing that they are not engaging in activity relevant to the tasks they are to learn when they are away from the testing situation. Thus, it is unlikely that animals will

rehearse task A in a retroaction design. It is possible, however, that human subjects will rehearse or otherwise engage in activity appropriate to the tasks they are required to learn in the experiment. Consequently, in most transfer experiments with human subjects, the control group is not kept completely inactive. Rather it is usually given some task to perform in the "rest" condition that keeps it occupied; this task is one that we know from previous work will have little or no effect upon the task in which we are testing for transfer. This is particularly important in studies of retention, for which rehearsal outside of the experimental learning situation would have serious consequences for the outcome of the experiment.

The Evaluation of Transfer

It is obvious that one evaluates the amount of transfer by comparing the performance of a control and experimental group of subjects. There are many different ways in which this comparison can be made, however, and sometimes they lead to quite different conclusions. Therefore the question of the evaluation of measures of transfer deserves some attention.

The ideal transfer measure would be one that could be expressed as a percentage in such a way that 0 per cent means no transfer and 100 per cent means perfect transfer, or the same performance level that would have been achieved if an equivalent amount of practice had been devoted to the task on which testing occurs. Gagne, Foster, and Crowley (1948) have given a formula which meets these requirements. It is

$$\frac{(\text{Score, transfer group}) - (\text{score, control group})}{(\text{Total possible score}) - (\text{score, control group})} \times 100$$

This formula is particularly useful in the proaction design in which there is no fore measure on the task receiving the transfer to take into account. By this formula, if the score of the transfer group is the same as the control group (which had not received practice on another task) there is zero transfer. If, however, the transfer group achieves the maximum possible score and if the control group does not (as would certainly be the case), transfer is 100 per cent.

This formula cannot be applied in every situation, however. In many examples of learning we do not really know what the total possible score is, or it is so unrealistic that we may end up with a distorted picture of the amount of transfer obtained. Furthermore, in the retroaction design we may want to take account of the effects of preliminary practice on the task receiving transfer, and this formula ignores such effects.

Expressing amount of transfer in such a formula is convenient, and it is

frequently done. It may however obscure the way in which transfer from one task to another occurs. It is quite possible, for example, that transfer may affect initial level of performance on a task but not rate of learning on that task (see Woodworth and Schlosberg, 1954). It is even possible that transfer could affect the upper limit of performance without affecting rate of learning, though this is unlikely. The point is, however, that expressing transfer in terms of some simple formula like the one above would obscure such differential effects of transfer. The particular parameter of the learning function that transfer affects may be very important in the practical evaluation of transfer in a training program (Lawrence, 1954), and we shall have more to say about this matter later.

THE THEORY OF TRANSFER

The Educational Backgrounds of the Problem of Transfer

Educational theories which concern transfer of training are older than the experimental study of the learning process. One does not need to look far to see why transfer of training is something that must be considered in any theory of education. In the first place, as we have already mentioned, education assumes transfer, for aside from its intrinsic value, there is no point to education apart from transfer. Furthermore, when the empirical study of the learning of children began, it was apparent that there were many practical questions about the transfer of training that demanded solution. Consequently, early educational psychologists were much interested in transfer; indeed, one could say that transfer of training was the central problem of early educational psychology. Since many of the theories and controversies that arose about transfer of training among the early educational psychologists are still with us (though usually in more artful and sophisticated guises), an examination of these will set the stage for the analytic study of transfer.

The notion of formal discipline. The idea of formal discipline, for a long time, was the *bête noire* of educational theorists. This was the idea, informally held by many people, that mental exercise is good for the mind. The classical preparatory curriculum in education embodies this notion. In former days the hapless student studied Greek, rhetoric, geometry, and Latin, not so much for their intrinsic value but because they exercised and sharpened the "mind." Learning to conjugate in Latin would make a keen wit as a lawyer, or training in logic would enable the physician to exercise the healing arts more effectively.

Despite the apparent embodiment of the notion of formal discipline in older educational programs, it never really had the support of thoughtful men (see Stroud, 1940). The evils of formal discipline, however, became a sort of straw man in educational theory, simply because 40 or 50 years ago the educators needed a good lever to pry the traditional curriculum loose from the schools.

Early studies of transfer. Early experimental studies of transfer of training were in large part directed towards demonstrating the inadequacy of the formal discipline theory of transfer—the notion that there would be transfer between two tasks to the extent that the first provided good hard mental exercise. Consequently there was much early emphasis upon the small amounts of transfer that actually took place between the formal academic disciplines. For example, Thorndike (1923) showed that there was only a slight advantage in the reading of English for students who had previously studied Latin. Similar studies (Thorndike and Ruger, 1923; Wilcox, 1917; Cole, 1924) showed the limited transfer of Latin to such varied aspects of English as reading, grammar, and spelling, as well as to the acquisition of modern foreign languages.

Relatively early, Thorndike and Woodworth (1901) performed a series of experiments designed to test theoretically some of the implications of the notion of formal discipline. In addition to experiments on classroom learning, Thorndike and Woodworth performed studies of transfer in laboratory tasks. For example, they practiced their subjects in the estimation of the areas of various geometrical figures such as triangles and rectangles. They examined the influence of transfer from learning to estimate the area of certain specific forms to the ability to estimate the area of other forms. The results of this experiment led Thorndike to an “identical-elements” theory of transfer. This notion and some that grew up in opposition to it are important, for they are the direct ancestors of our current interpretations of the mechanisms in transfer.

Identical elements versus general principles in transfer. The identical-elements theory of transfer says that the specific elements identical to two tasks are transferred from one task to the other. Thus, training in addition benefits performance in multiplication for the simple reason that many of the responses required in addition are the same as those required in multiplication. From the notion of identical elements, then, one would expect the amount of transfer to be determined by the number and importance (relative frequency of occurrence) of responses common to two tasks.

This concept is a vast improvement over the omnipresent transfer pre-

dicted by such a notion as formal discipline, for the theory of identical elements clearly predicts that transfer will differ in amount between two tasks according to the variations in the tasks themselves.

Educational theorists, however, have not always liked the notion of identical elements because of its suggestion of an atomistic structure in what individuals learn. People learn, they say, not just isolated responses which may work in some new situations and not in others; they learn general principles and rules for applying these principles in a rational way. Thus, to use the same example we used to illustrate identical elements, it is not merely that a child can apply what he has learned about sums to multiplication of n -digit numbers; it is that some general principles about the properties of decimal numbers apply to multiplication. It is as if a bright child were taught the basic principles of addition and subtraction, and from these, himself discovered multiplication and division.

For a while educational psychologists and learning theorists acted as if there were two mutually exclusive notions about transfer—one right and one wrong. Common experience, however, suggests that transfer can occur by both identical elements common to two tasks and general principles that govern two or more tasks. The basic psychological problem is not to decide which of these two possible modes of transfer is correct or more valuable, but to attempt to uncover the basic rules governing all kinds of transfer. This has been the goal of most learning theorists interested in transfer, and, in part at least, they have been successful. The most elementary step, of course, is a stimulus-response analysis of the transfer problem, and that is what we shall examine next.

The Stimulus-Response Analysis of Transfer

A stimulus-response analysis of transfer is important because the most elementary way two tasks can differ from one another is to differ in some simple stimulus or response component. Thus if we train an organism to make some response to a particular stimulus and then test for this response with a slightly different stimulus, we have an opportunity for the evaluation of transfer. We can see that this is the kind of situation in which we measure stimulus generalization; indeed, this generalization turns out to be a special case of transfer of training, perhaps the most elementary. Let us examine it as a special case of transfer.

Stimulus variation with response held constant. The essential characteristic of the measurement of stimulus generalization is that we hold some response constant and vary the stimulus that has either been conditioned to

the response or has been consistently paired with the reinforcement of that response. Remember that the major results of the study of stimulus generalization showed that the more similar the test stimulus was to the original training stimulus, the greater the potential for eliciting the original response. From these studies we should guess that the amount of positive transfer would gradually decline as the stimulus situation was changed from what it had been during original training (assuming the responses required remain the same). Such a relationship is illustrated in Figure 46. On one axis in this figure is the amount of (positive) transfer from 100 per cent to zero and on the other axis is stimulus similarity. Notice that 100 per cent transfer

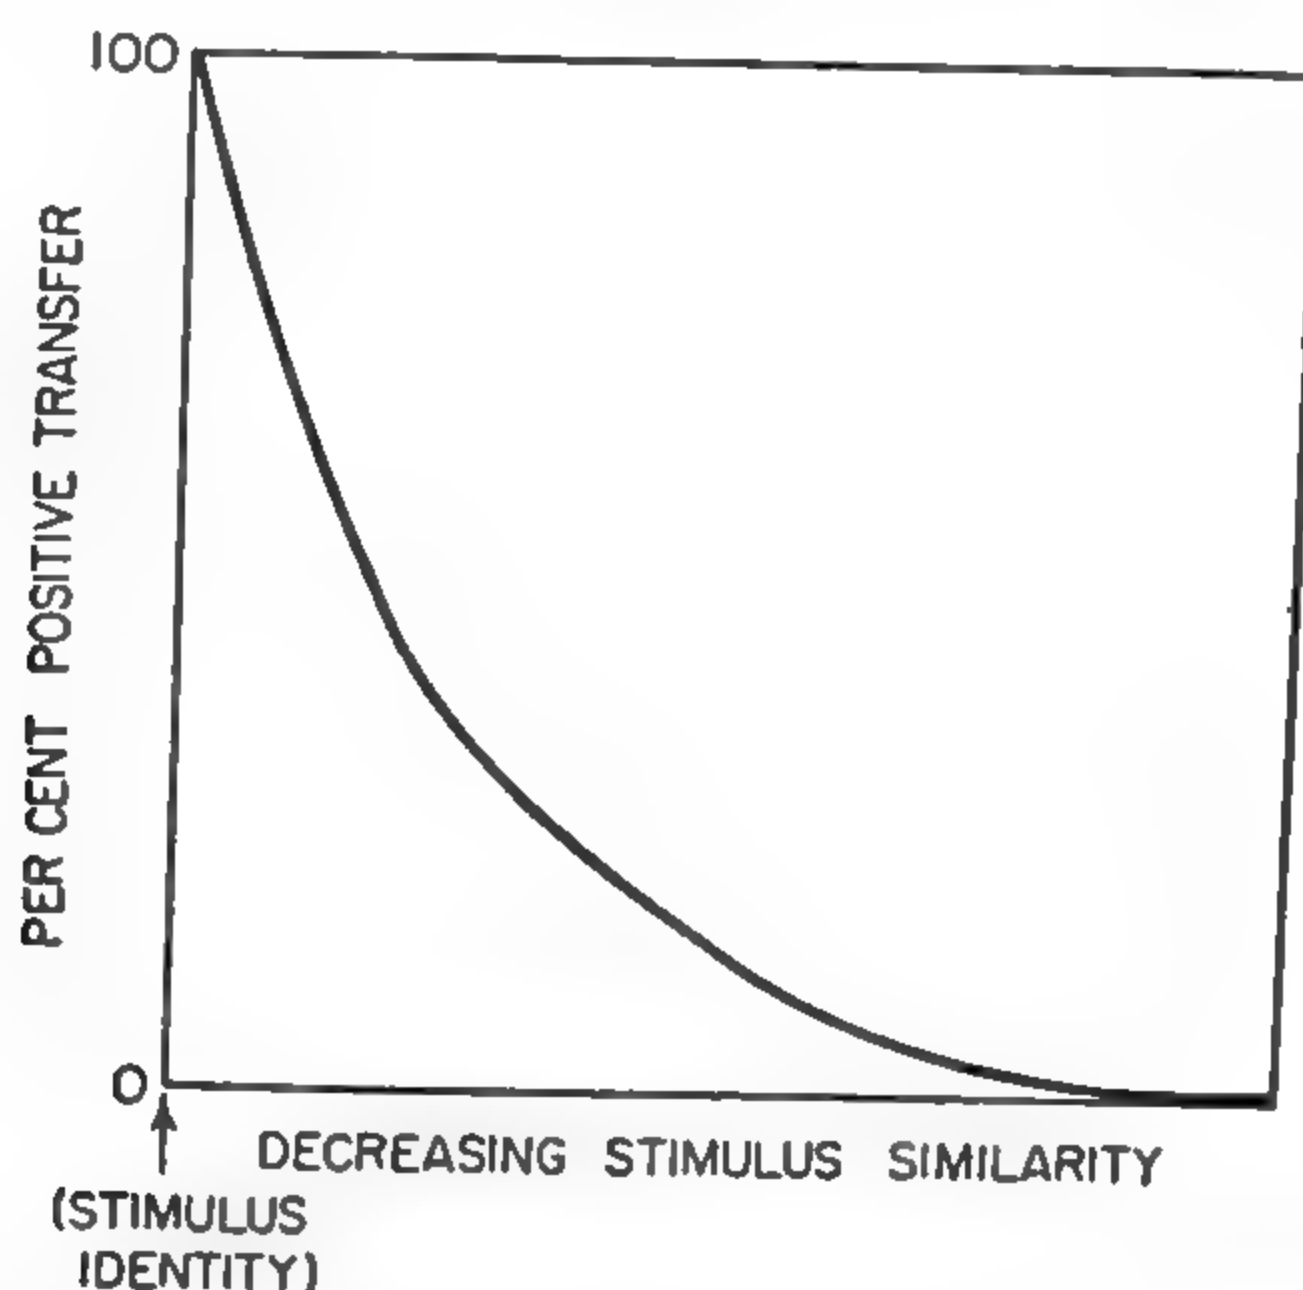


Figure 46. Per cent transfer of training between two tasks when the responses are identical and the stimuli are varied from being identical to being completely dissimilar.

is defined as that case in which the test stimulus is exactly the same as the one used in original training.

We may ask, Does the kind of relationship shown in Figure 46 hold for more complicated examples than conditioning? That is a difficult question to answer, since procedures used in the study of different learned acts are likely to be very different. Let us, however, examine some relatively simple examples of verbal learning.

In one well-known experiment (Yum, 1931) subjects learned to associate nonsense syllables with stimulus items. Each syllable was paired with a different item, and after the subjects had learned all the associations, they could recall the appropriate syllable when they were given the proper stimulus item. In some conditions the stimuli were other nonsense syllables and

in others they were abstract geometric forms. Let us look at one condition in which the stimuli were the latter. The subjects first learned to pair each form with a different nonsense syllable (there were 12 in all). Then, 24 hours later they were presented with the stimuli once more and asked to recall the right syllables. In some conditions the same stimuli as those used in original training were presented. In other conditions these stimuli were altered so that they varied all the way from resembling the original stimuli very closely to resembling them only vaguely. The results of the recall test were as follows:

<i>Stimuli</i>	<i>Per cent recalled</i>	<i>Per cent of recall to identical stimuli</i>
Identical	84.6	100
Highly similar	64.5	76
Similar	49.2	58
Less similar	45.3	53
Least similar	36.3	43

We can see that as the similarity of the stimulus items to those used in original training decreased, the percent of items recalled decreased. Thus, there was less positive transfer with dissimilar stimulus items than with similar ones.

Results much like these were also obtained by Gibson (1941), but her experiments went one step further. This is important because one could argue that the subjects in the experiment just described did not respond to the nonidentical forms as frequently as they did to the identical ones simply because they did not know that they were supposed to. Gibson, however, studied *relearning* of the same responses when the stimuli were identical or somewhat varied. The results were much the same; there were greater savings in relearning when stimuli were identical or highly similar than when they were only slightly similar or dissimilar.

Changing responses and holding stimuli constant. What happens when the stimuli for two different tasks are the same but when the responses are different? There is an old rule of thumb that says this situation always produces negative transfer. This implies that the second task is harder to learn than it would have been if there had been no first task. The further implication is that the responses learned in task one interfere with the acquisition of new responses to the same stimuli. We shall deal with this implication at length in the next chapter.

For now, we should note that this rule is a considerable simplification of what really happens in practice. Actually, it sometimes happens that

there is positive transfer when responses are changed and the stimuli are the same. Let us look at some examples.

In one experiment (Kellogg and Walker, 1938) dogs were conditioned to flex their right rear legs to a buzzer by pairing the buzzer with a shock to that leg. The shock was then transferred to the left rear leg. The animals learned this response much faster than they did the original right-leg response. Further exploration showed that it was easier to recondition the right leg than to condition the left leg after original conditioning. Furthermore, it turns out that there was a rough gradient among the four feet, so that after the conditioning of the right rear leg, the left rear leg was easiest to condition, then the right front leg, and finally the left front leg (Kellogg and Wolf, 1940).

In this situation the transfer is evidently a matter of components of the responses in original training being common to the transfer situation. In the initial conditioning of the right rear leg the animals sometimes respond with the other feet as well. The degree to which the conditioning of the right rear leg transfers to the other feet depends upon the probability that the animal will respond during original conditioning with one or more of the other feet (it should be mentioned that in these experiments the dogs were supported in such a way that they did not need to use their feet for support).

What about the case in which changing the response leads to negative transfer? This would happen if a subject learns the association S_1-R_1 and then finds it more difficult to learn the association S_1-R_2 than if he had not made the original association. In one of the best-known experiments on this problem (Bruce, 1933) the associations consisted of pairs of nonsense syllables. Thus the subjects would be required to give a response like KIV when given the stimulus REQ. For the second list the subject might be required to give a response like ZAM to the same stimulus REQ. Under these conditions negative transfer seems to occur. Learning a new response to an old stimulus produces interference.

Not all experiments, however, have found this to be so. For example, a more recent experiment (Porter and Duncan, 1953) which used the same design, except that two-syllable adjectives were used instead of nonsense syllables, showed no negative transfer when a new list of response items was learned to old stimulus items. The lack of negative transfer was partly explained by the fact that since the subjects in this experiment were not highly practiced at rote verbal learning, simply learning some of the techniques of this kind of memorization obscured some of the negative transfer. There is, however, a more fundamental reason why negative transfer

did not occur, for in this experiment negative transfer was found with slightly different conditions. If the second task consists not of an entirely new list of response terms learned to the same stimuli but of the same response terms shuffled so that they are paired with different stimulus items, negative transfer occurs. Thus under this condition negative transfer was not swamped by general-practice effect. The reason, of course, was that the interference was much greater because the responses in the second task were already connected to other stimuli in the list making up the first task.

Thus, in general, it appears that if all other things are equal (this usually means that there is no practice effect getting in the way), responses which are mutually exclusive will interfere with one another when we try to attach them to the same stimuli. In the experiment we mentioned a few paragraphs ago on the conditioning of different legs in dogs, the responses were not mutually exclusive or antagonistic, and they did not interfere with one another; indeed, since they were all part of common components of reaction to electric shock they tended to facilitate one another.

This difference of positive transfer between responses learned to the same stimuli when they are components of the same act and negative transfer when the responses are mutually exclusive is sometimes expressed by saying that the kind and amount of transfer depends upon the amount of response similarity. Response similarity means that when a response is associated to a given stimulus there are other responses which will tend to be elicited by that stimulus. It is assumed that there exists some gradient of response similarity such that, to a given stimulus, some responses have a greater probability of being elicited than others. We can see that the notion of response similarity is roughly analogous to stimulus similarity.

The transfer surface. We can put all of the stimulus-response relationships in transfer together in one diagram, but the diagram will require three dimensions, one for stimulus similarity, one for response similarity, and one for amount and direction of transfer. Such a diagram, derived from one originally drawn by Osgood (1949), is shown in Figure 47. The part of the figure in the lower right-hand corner is where maximum negative transfer is found. This is the case in which the stimuli between two tasks are the same but the responses required are different. On the other hand, maximum positive transfer is found in the upper left-hand corner where stimuli and responses are identical—in other words when the second task is merely further practice on the first. From this diagram we can see that further practice on the same task is only a special and limiting case of positive transfer, and perhaps logically it should be. On the opposite side of the dia-

gram, which is the side on which stimuli are unrelated, transfer is very weak. The limiting case here is the extreme right end of the diagram. Here the stimuli between the two tasks are unrelated and there is no similarity of responses. Therefore, no transfer, positive or negative, will occur between the two tasks. We should not expect learning how to swim to have much transfer to performance at the typewriter.

This diagram summarizes all of the possible stimulus-response relationships between two tasks. It ought not, however, be taken too literally. For

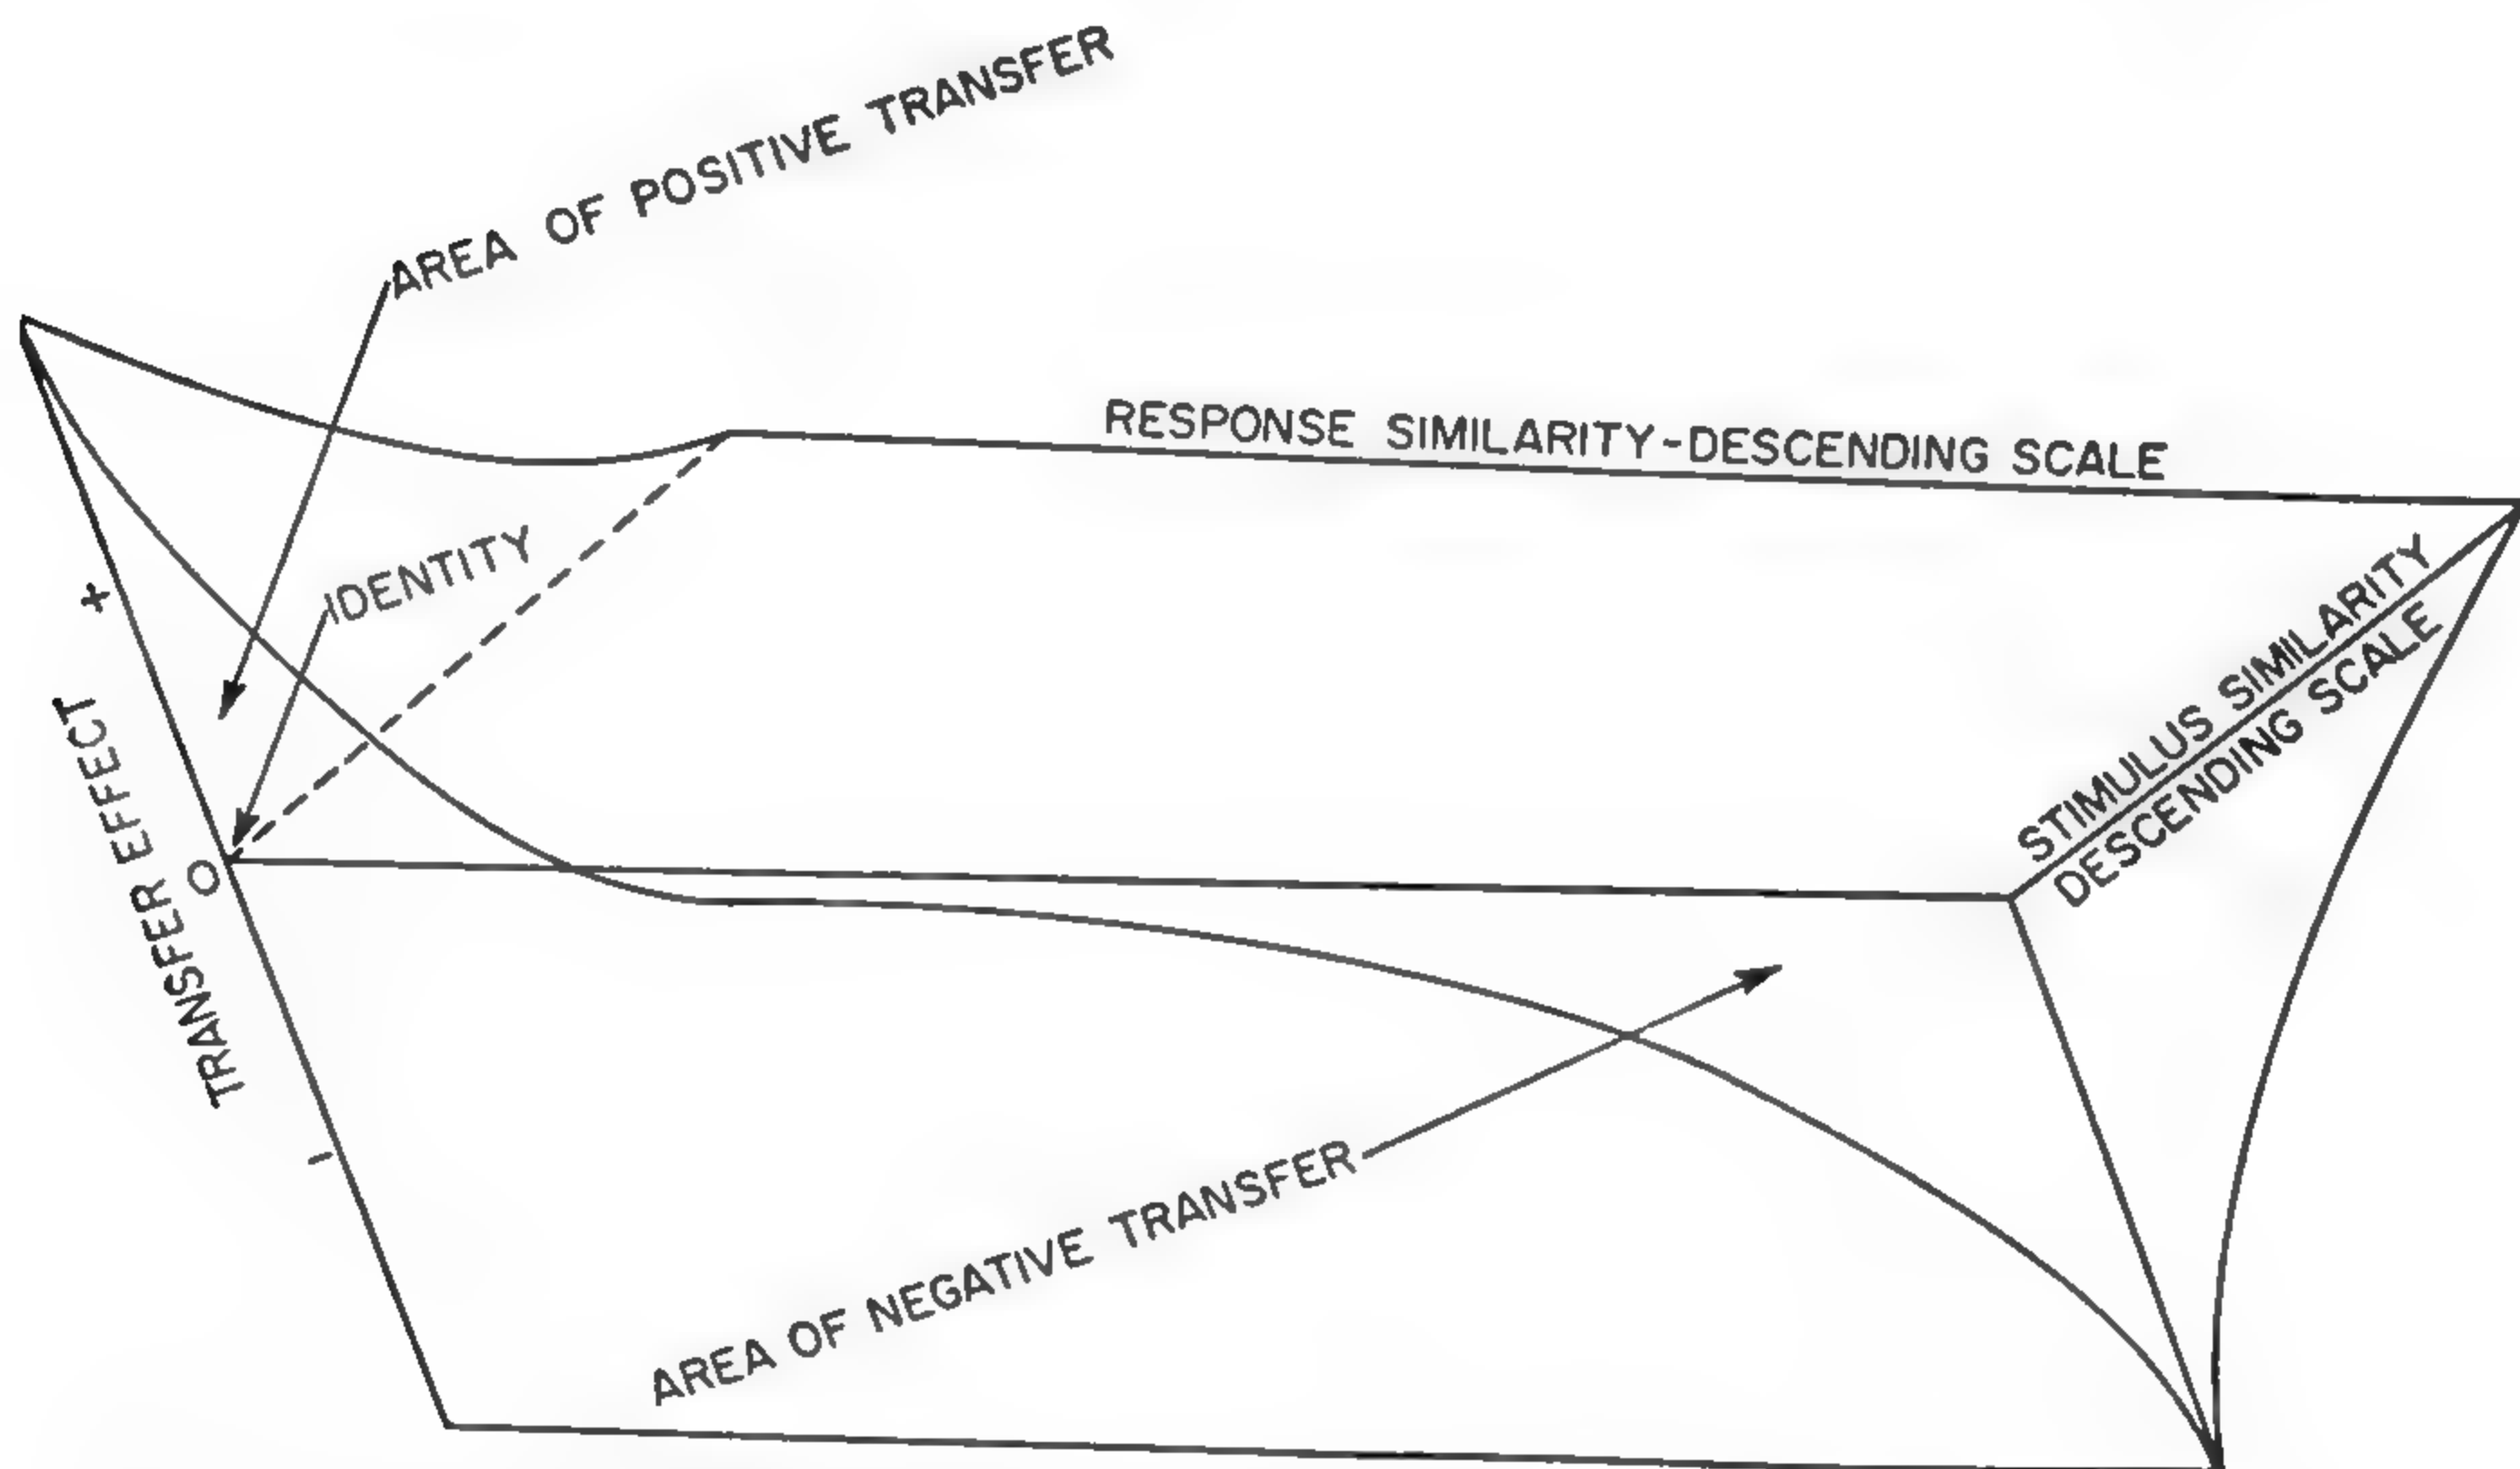


Figure 47. The hypothetical transfer surface. The amount of transfer (positive and negative) as a joint function of stimulus similarity and response similarity. (After Osgood, 1949.)

one thing—as we shall shortly see—there are apparent exceptions to the rules implied by the diagram. For another thing we do not have general quantified scales for stimulus and response similarity. Therefore the exact shape of the diagram has no meaning; only the general relationships are meaningful. If any of the dimensions were expanded or contracted at the expense of any others, it would not do violence to the basic notion presented in the diagram. In short, it is a qualitative picture of the stimulus-response relationships in transfer, and with the exceptions noted below, it works fairly well in summarizing data on transfer.

Transfer situations not summarized by the transfer surface. There are many kinds of transfer situations that are not easily summarized by the

transfer surface. Among these is learning how to learn, the general improvement that comes in learning many different tasks all of which demand a common technique of learning. In theory, it can be argued, this kind of transfer can be considered as a special case of the transfer surface, by thinking of the technique of learning (as it occurs, say, in rote verbal memorization) as consisting of a series of highly similar responses involving mode of attack, etc., which run through all the tasks to be learned. We are not always sure what such responses would be, however, and considering the general problem of learning how to learn as a special case of response similarity leads more to the aesthetic satisfaction of having one simple principle of transfer than to a real understanding of the mechanisms in learning how to learn.

More seriously, there are real exceptions to the rules implied in the transfer surface. For example, Young and Underwood (1954) note that positive transfer does not always occur when stimuli are different and responses are similar to one another. They showed that there will be little relationship between response similarity and amount of positive transfer if the responses *within* the two tasks are not difficult to differentiate. In other words, similarity of response between two tasks will not be an important variable in transfer unless the responses within the tasks are difficult to differentiate one from another. Thus the extent to which the transfer surface applies depends upon the stimulus-response relationships within tasks.

An even more serious difficulty with an uncritical use of the transfer surface is indicated in an experiment by Deese and Hardman (1954). In this experiment subjects learned to complete sentences by the method of paired associates. Thus a statement like "the path is" would be given as a stimulus item, and the subjects would learn to complete the statement by giving the correct word "rocky." A series of such statements were given as the original task. In the second task, the same stimulus items were given but the subjects had to learn new response terms. In one case the new response terms were synonyms (thus in the case of "the path is," the new response would be "stony"); in the other case the new response terms were unrelated to the old ones (thus to the example given, the new response would be "short"). Looking at the transfer surface in Figure 47 we should probably expect the synonyms to yield positive transfer and the unrelated words to yield negative transfer. And that is exactly what happened in the learning of the second task. So far so good. This experiment went one step further however; the subjects were asked to relearn the original list. In this relearning, there was negative transfer in both cases and the amount of

negative transfer was greater after learning the synonyms than after learning the unrelated words. Thus, in this experiment, positive transfer was found in proaction with synonyms and negative transfer in proaction with unrelated words. In retroaction, however, both cases produced negative transfer, greater in the case of the synonyms, in apparent contradiction to the transfer surface.

The explanation for this finding is that the synonyms in the second task were easier to learn and hence better learned in a fixed number of trials than the unrelated words. Thus they competed more with the retention of the original items than did the relatively weaker unrelated words. This is a reasonable explanation that does not violate the spirit of the transfer surface, but it should make us cautious in making a blanket application of the relations in the transfer surface. In general it is less easy to go astray in applying the transfer surface to proaction than to retroaction experiments.

Identical elements and common principles in the transfer surface. The transfer surface is a lineal descendant of the theory of identical elements in transfer. This is because most of the experimental studies of transfer have emphasized the analytic aspect of the problem, and the relationships between tasks have been dissected into their stimulus and response components. The similarity relationships that have been studied can easily be considered in terms of common elements. Therefore the transfer surface can be taken as a kind of detailed account of the transfer one would expect with identical elements in two tasks.

What about transfer by general principles? It is not difficult to stretch the same stimulus and response analysis to this case. One well-known study of transfer in a schoolroom task illustrates this. Knight (1924) had two groups of children practice the addition of fractions. One group practiced on fractions with denominators ranging from 2 through 30. Another group practiced with a limited number of denominators, 2, 4, 6, 8, 12, and 24. A test was given on those denominators practiced by the first group but not by the second. Both groups did just about equally well on the test. In other words transfer was perfect. It is easy to see that this example, which fits the description of transfer by general principle, also is in accord with the transfer surface. In this case some component, a general rule, was common to the two tasks.

Wherever possible we shall try to refer examples of transfer of training, both in this chapter and succeeding chapters, to the transfer surface. While some experimental situations do not lend themselves easily to a simple

stimulus-response analysis, the general usefulness of the transfer surface is considerable.

SPECIAL PROBLEMS IN TRANSFER

Learning How to Learn

One of the most general and important kinds of positive transfer is *learning how to learn*. Every investigator of verbal learning knows that subjects who have never learned by rote association do relatively poorly in laboratory studies of paired-associate or serial-anticipation learning. Such subjects will take a large number of trials to learn their first list of, say, nonsense syllables. They generally show great improvement on the next list, however, and on each succeeding one they show additional improvement. After 15 or 20 lists, however, the amount of improvement from list to list is relatively small; the subjects have learned just about all they can about the techniques of learning nonsense syllables by rote.

Learning how to learn is a very general phenomenon. So far as we know, it occurs in some degree for all techniques of learning and for all kinds of tasks. One of the central problems in the transfer of training is to find out how and why such improvement occurs, as well as to see what the implications are for more general problems in the psychology of learning.

Learning sets. In the chapter on discrimination learning we examined one such attempt to understand the mechanisms in learning how to learn. This particular problem is the formation of learning sets and is specifically concerned with the extent to which improvement occurs in successive discrimination problems.

We have seen in our discussion of discrimination learning that as organisms are given successive problems, their ability to solve each problem improves. Thus it may take a monkey many trials to learn that a raisin is to be found under a red box instead of a blue cylinder if this is the first discrimination test it has ever worked at. After a number of such tests, however, the monkey needs only one trial to show perfect discrimination.

An analysis of the learning-set problem showed that it was the result of reduction of stimulus generalization (and hence negative transfer, since each problem requires a different stimulus-response connection) and the gradual emergence of a new habit, a habit of testing stimulus objects irrespective of every dimension except position (left or right) for their correlation with reinforcement. The implication is that through the reduction of stimulus generalization and negative transfer a new habit is allowed to

emerge which transcends the particular problem the animal is working on. The animal learns to look for the significant variation in stimuli, and his solution is likely to appear to be insightful.

The development of a learning set can be characterized as the emergence of a higher-order habit composed of the single common component reinforced in a number of lower-order habits. Each discrimination problem can be looked upon as a lower-order habit which must be mastered. The only behavioral feature that is consistently reinforced from problem to problem is the habit of making a choice and changing the choice when the initial one is incorrect. The habit of choosing blue stimuli because of correlation with reward in one problem or square stimuli for another problem will only be reinforced part of the time. The only aspect of the set of discriminations consistently reinforced is that of reversing the choice of stimuli when the initial one is incorrect and keeping with the initial choice when it is correct.

Do learning sets occur outside discrimination-learning problems? As we shall see later, there is an analogy between learning abstract concepts and learning sets. Furthermore, probably many examples of learning how to learn are the result of emergence of higher-order habits. Moreover, even the relatively simple examples of learning how to learn illustrated by discrimination learning in monkeys involve such diverse factors as learning not to perseverate to the same stimulus from trial to trial and learning to respond to differential aspects of the cues (Harlow, 1950b).

Teaching how to learn. We have seen that there is a general improvement in ability to learn that comes with practice at a particular kind of learning. There is more to it than this, however. Many kinds of remedial educational programs are built on this assumption, plus the supposition that special methods of teaching how to learn will produce even greater improvement.

In a classical experimental study Woodrow (1927) showed that specific instruction in how to memorize poetry, prose, and factual material resulted in greater improvement in ability to memorize than undirected practice. Woodrow stressed many of the techniques that are now standard in teaching students how to learn. He used active recitation, attention to meaning, techniques for maintaining continuous alertness, etc. Woodrow's study was not analytical, and we cannot tell which of the factors was most important in improving memory. The important point, however, is that training for memory does have beneficial effects over and above those that come from practice alone.

Differentiation and Transfer

In the section devoted to stimulus-response analysis and theory of transfer we pointed out that one important interpretation of stimulus and response relationships was that transfer was to a considerable extent mediated by differentiation. This was the view of Gibson (1941), and it has received considerable experimental attention (Young and Underwood, 1954; Underwood, 1954a). One of the implications of this notion is that if we could induce differentiation among either the stimuli in a task or the responses *without specific practice on the task or on the responses in it*, we could produce positive transfer.

A simple example of stimulus predifferentiation is illustrated in an experiment by Gagne and Baker (1950). These investigators gave subjects the task of moving a switch when a particular light came on. There were a number of lights located in different positions, and each light signaled a separate switch. Some of the subjects were given preliminary practice, not in moving the switches, but in merely assigning letters of the alphabet to each of the lights. This practice in associating the lights with names had an enormous influence on learning to associate the lights with the switches, even though the switches had nothing to do with the letter names. Those subjects that had no stimulus-predifferentiation training made about 7 mistakes in their first 20 trials on the switch task. Those who had 32 stimulus-letter naming trials made only about 3 errors.

This experiment was interpreted to mean that stimulus predifferentiation reduces the amount of stimulus generalization between the various stimuli so that in the final task, when these stimuli are associated with motor responses, they are more distinctively different for the subjects.

There have been many similar experiments demonstrating essentially the same thing. Goss (1953), however, has pointed out that it is not entirely certain that the obtained transfer in these studies has been due to stimulus predifferentiation alone. He has further analyzed this problem and shown that while stimulus predifferentiation is indeed the main factor in the transfer effect, other things, such as warm-up, also contributed. In general, any kind of labeling or differentiation of stimuli will help in any subsequent task that demands different responses be attached to these stimuli.

Transfer of verbal learning to the acquisition of a motor skill is common, and it very likely occurs principally because of the differentiation which preliminary verbal training allows. Battig (1956) reminds us that the learning of motor skills involves both discovery and performance. Discovery

means learning to differentiate between stimuli, to pick out their critical aspects, and to learn which ones are to be reacted to in a different manner. Performance means the actual motor responses which must be made. In many motor skills preliminary verbal training will suffice for learning the stimulus relationships. Battig showed that the amount of transfer from verbal pretraining to a motor skill varied inversely with the complexity of the motor activity required. When the motor, or response, side of the skill is relatively unimportant, verbal pretraining is of great value. If the motor aspect is too complicated, there is little or no transfer to the motor skill. If these results are generally true for many different tasks, we may have an important principle for the application of verbal pretraining to learning skilled performance; the value of such training decreases as the motor complexity of the task increases.

Transfer and Difficulty

There are some rather surprising and complicated relationships between transfer and the relative difficulty of the tasks involved. In some cases transfer from an easy task to a hard one is greater, and sometimes the converse is true. Let us examine the latter case first.

An experiment by Jones and Bilodeau (1952) illustrates the situation in which transfer is greater from a hard task to an easy one. Both tasks these investigators used involved coordinating the two hands to keep an indicator on a moving target. The difference between the two tasks was in the complexity of the movement of the target. In the hard task, the target moved in a very complicated path; in the easy one it moved in a relatively simple path. With the use of several groups of subjects, the investigators studied the transfer from the hard path to the easy one and vice versa. There was very great transfer from the complicated path to the easy one; as a matter of fact the transfer in this case was equivalent to *direct* practice on the easy task. There was also positive transfer from the easy task to the complicated one, but not nearly as much.

In the particular situation studied by these investigators, the harder task involved all of the components of the easy task and permitted greater practice. The easy task had only some of the same response components as the difficult task. Therefore the transfer, while positive, was somewhat less. In general, this will be true; in tasks in which more difficult versions contain all of the stimulus-response elements in simpler versions plus additional ones, positive transfer may be large, indeed, as large as the effects of direct practice in instances where no negative transfer from particular stimulus-response elements occurs.

In other situations, however, it seems that preliminary practice on easy tasks produces greater transfer than preliminary practice on difficult material. This was the case in a study of discrimination learning in rats (Lawrence, 1952). The rats were trained to discriminate between the brightness levels of two goal boxes. In one condition the animals received all of their training on a difficult problem, that is to say, a problem in which the brightness levels were nearly the same. In other conditions the animals received preliminary training on problems in which the brightness levels of the two boxes were quite different, so that discrimination was easy. All animals were eventually tested on the difficult discrimination.

The results, rather surprisingly, were that all the animals trained on easier problems did better than the animals trained on the difficult problem. This result is probably the result of factors much like those responsible for the stimulus-predifferentiation effect. Training on the easier discrimination allowed the animals to discover more readily the critical difference between the two boxes and pay attention to it and to no other stimulus conditions. Thus, when these animals were transferred to the more difficult discrimination they were alerted to brightness differences and found it easy to solve the problem. Animals trained all along on the difficult discrimination, however, found it difficult to solve the problem because the small brightness difference between the two boxes was never brought directly to their attention.

Thus in cases in which difficult problems involve stimulus-response components identical to those in easier tasks (but also additional stimulus-response components) there will be more transfer from the difficult problem than from the easier one. In those cases in which the problems involve the isolation of relevant stimulus dimensions and the association of certain stimuli with reward or solution of a problem, training on easy problems will transfer more than training on difficult problems.

These two different cases have not been completely explored, but it is likely that the principles in them can be generalized to many different tasks. If so, we have some important rules about the application of transfer of training to problems of varying difficulty.

POSITIVE TRANSFER IN SCHOOL LEARNING

We have already seen that transfer of training has been one of the problems of major concern to educational psychologists. We shall not look at any of the practical problems of the applications of transfer principles to the school situation, but we shall look at some of the more general ques-

tions that have come from the study of transfer in school. One of the most important of these is training for transfer.

Training for Transfer

One of the basic problems in teaching is to maximize the amount of transfer. We have seen that transfer from one task to another may not be very large; consequently educational psychologists have looked for ways to increase the amount of transfer. Mostly this has meant that psychologists working in the school situation have tried out various methods for making transfer greater between academic subjects, or an academic situation, and some practical problem. An example will show the sort of things educators have been concerned about in transfer.

An example of training for transfer. It is generally accepted that a formal academic discipline like high school geometry provides some transfer to reflective thinking in general. In some circles this is regarded as a hold-over from the notion of formal discipline, but, of course, such transfer does occur. The postulational technique demonstrated in geometry does generalize to other things, and students are sometimes able to see this. The late Clark Hull, as a matter of fact, said in his autobiography that his interest in scientific theory generally came from his high school days and the study of geometry.

In the main, however, transfer from a traditional geometry course to other kinds of situations involving reasoning is disappointingly small. Ulmer (1939) performed an experiment in an effort to see what could be done to improve the amount of transfer by designing a special course in high school geometry for transfer.

Ulmer used three large groups of high school students in his study. One group served as a control. A second group of students was taught geometry by the traditional methods. A third group was taught geometry through a program designed to arouse critical thinking both connected and unconnected with geometry. The teachers for the students in this group used a special manual designed to extract as much training in general critical thinking from geometry as possible. The results of this study were very clear-cut. On tests of syllogistic reasoning and in the detailed records of teachers' log books it was clear that the students in the third group were able to transfer much more adequately the basic principles of logical thought as exemplified in geometry to the tests of reasoning than the group trained in the traditional method.

What are the conditions of training for transfer? Ulmer's study gives

TRANSFER OF TRAINING

impressive evidence for the possibilities of training for transfer, but it does not tell us about the nature of training for transfer. We cannot know, from the available evidence, the complete answer to this question, but we do know enough to be able to see some of the more important factors.

One paramount factor is the control of generalization. We saw in studies of animal learning and human rote verbal learning that generalization is an essential element in both positive and negative transfer. We also saw in the experiments on the formation of learning sets that experience at multiple-task learning gives learners direct control over generalization. That is to say, organisms learn to *vary generalization voluntarily*. Variation in generalization is learned as a critical element in problem-solving behavior, and the sophisticated learner (be it rat or man) varies the extent to which behavior is generalized in accordance with other things learned about the testing situation. Thus the monkeys in the learning-set experiments learned to test specific hypotheses about which stimuli would lead to reward, rather than to rely simply on the strongest generalization of the moment.

This is extremely important, and we shall discuss it again in connection with problem solving. It should be clear in the present context, however, that broad generalization is essential for transfer. Unfortunately transfer is both positive and negative, depending on whether the generalized behavior is compatible with the requirements of a new situation or not. A higher-order behavior is achieved when the organism can decide whether or not to vary transfer based on its estimates of the relative probabilities and advantages of negative and positive transfer.

There are numerous applications of particular techniques of training for transfer. Many can be found in basic textbooks in educational psychology (for example, see Frandsen, 1957). Such matters as unifying concepts in arithmetic and phonetics in spelling are all important questions of transfer in the elementary curriculum. Indeed, a good deal of the dispute over the use of the so-called "phonic" method in reading is basically concerned with whether this method produces mostly negative or mostly positive transfer for reading at higher levels. The opponents of the widespread use of this method have emphasized the potential sources of negative transfer. The proponents of the method have emphasized its wide generality, and hence by inference, its high potential for positive transfer. Since almost any rule or principle in language learning can potentially produce both negative and positive transfer, much of the concern over different methods of teaching language skills involves the question of transfer.

To emphasize the potentiality for negative transfer in language learning,

consider some of the erroneous definitions given by university students in a vocabulary test (Johns, 1939):

Elucidation

To get away from ("elude")

Elimination (from the sound)

Coerce

Speak, discuss (possibly from "converse")

Plan of study (from "course")

Innate

Without life, dead (from "inanimate")

Silly ("inane")

Obviously, some training in the fundamentals of English words would provide enough control over generalization so that positive rather than negative transfer would result.

Ability and Transfer

In the study of Ulmer on the teaching of geometry, students of all levels of intellectual ability seemed to profit from the specialized methods of teaching designed to increase transfer. There was, however, a relationship between the gains from transfer and intelligence. This relationship held both for the experimental condition, in which the specialized method of teaching was used, and for the traditional method of teaching geometry. There was always more transfer from geometry to improvement in syllogistic reasoning for intelligent students than there was for average or inferior students.

This relationship between transfer and intelligence is characteristic of many studies. As a matter of fact, it almost looks as though the ability to recognize and make use of a general principle when it is available is a basic factor in intelligence. Most of the educational implications of such a relationship are quite obvious. There is one important implication that is perhaps not so immediately easy to see. Transfer by general principle may not work well with younger children because of the correlation between age and mental maturity and the dependence of transfer by general principle on mental maturity. In one study of second grade children, for example, it was found that simply showing the children that arithmetic procedures could be generalized was more productive of transfer than attempting to teach them rational procedures (Overman, 1930-1931). These second grade children were probably not mature enough to grasp the general principles involved. They had to be taught, more or less by rote, to apply the

rules that they had learned. As we might suppose, those children higher in mental ability profited more from the rationalization procedure.

The relationship between transfer and mental maturity plays a considerable role in an important educational concept, the concept of readiness. Educational psychologists have made much of the futility and, indeed, danger inherent in trying to teach children skills before they are ready for them. In general this means that one should not try to enforce practice at a skill until certain basic functions have matured to a certain level. It is impossible to teach most preschool children to write, because they have not yet the command over the fine muscular movements that are required in writing. Such children can only make rather large and clumsy movements with their fingers and arms. It is to be noted, however, that a considerable (though indeterminate) part of the maturity required for the ability to write comes from previous training. Scribbling, drawing, working with the paint brush, and indeed, a whole host of related activities probably all give the child an ability to work with the muscles of his hands and arms that later transfers to writing.

Transfer and Specialized Training Programs

The value of general education can exist quite apart from any specific transfer value assigned to parts of it. This is not true of most specialized training programs, however. If the Army or the Air Force invests time and money in training a man for a particular job, it wants to know whether this training has actually made him better for the job than he was before training. In other words, in specialized training programs transfer is a critical element.

In many applied problems in transfer the important problem is that of validity. We need to ask the question, Does training on a special device or by special techniques show positive transfer to the actual task? There are many questions subsidiary to this main question. For example, on-the-job training is usually very expensive. Will a preliminary training program not only transfer, but will it transfer enough to the actual task so that the training program is, in the long run, cheaper than on-the-job training? Such problems are of very real concern to the applied psychology of transfer, and they have received much attention (Lawrence, 1954; Tufts Institute, 1952). It is too early to hope that the basic experimental work on transfer will allow us to predict well enough to answer these questions without a direct test in the applied situation, but it is not too much to hope that we shall be able to make fruitful guesses about which applied training programs will profit most from a revamping in view of the basic principles of transfer.

CHAPTER 10

RETENTION AND FORGETTING

There are two ways of looking at the problems of learning. We can examine the way in which behavior changes as a function of experience; ordinarily this means that we view improvement in performance as the result of practice. We can also take a fixed level of performance and see how that performance is retained over an intervening time interval. This is the method we shall consider in this chapter. Both of these approaches are, of course, simply different views of the same problem. If there were no retention, there could be no learning, since the gradual increase, or improvement, that occurs with practice is based upon residual benefits from previous practice. This suggests that the factors underlying retention are the same as those underlying learning and, conversely, that the factors underlying forgetting are the same as those which tend to slow down the learning process. If this is so, why a special study of retention?

The answer to this question is that such a study allows us to examine changes in behavior uncontaminated by further practice or exposure to learning materials. When we have determined to what extent and in what manner something is retained, we have a measure that tells us how mental processes reorganize material already learned without additional exposure to the material. From a study of this process of reorganization we can learn something about the mental processes themselves.

Before we begin an examination of retention, we need to define the two words retention and forgetting. *Retention* refers to the extent to which material originally learned is still retained, and *forgetting* refers to the portion lost. Therefore

$$\text{Amount forgotten} = \text{amount learned} - \text{amount retained}$$

Amount retained is always measured directly, and amount forgotten is obtained by subtraction. Since amount retained is the direct measure, it is the measure we prefer to use. It is useful, however, to talk about forgetting when we are considering those conditions which tend to reduce the amount retained, and this chapter is therefore separated into two sections—one called Retention and one called Mechanisms of Forgetting. In the second section, emphasis will be on forgetting as an active process, and we shall discuss those experiments in which there is a deliberate attempt to reduce the amount of material people will be able to retain. In the section on retention we shall examine studies of the amount and the type of things people remember after various intervals of time.

RETENTION

Retention was one of the first aspects of human learning and verbal behavior to be studied by laboratory methods. Ebbinghaus studied retention as part of a sustained attack on the “higher mental processes.” The results of his studies were published in a monograph in 1885 called *Über das Gedächtnis*, or *Concerning Memory*. This monograph is a landmark in experimental psychology, not only because it reports important findings but because it represents the first experimental study of the problems of human verbal behavior. This opened up the whole field of experimental investigation of thought, memory, and verbal learning.

One of the problems which Ebbinghaus studied was how much can be remembered after a certain time interval has elapsed since learning. It is rather surprising that no one had ever thought to do this before, even in a casual way, but the fact is that Ebbinghaus’s results represent the first really quantitative data on human memory. Before we look at his findings, however, and those of his successors, it will be wise to examine some of the methodological problems in the study of retention.

The Measurement of Retention

Experimental Designs

Experimental methods in the study of retention fall into two classes. In one class the same subjects are tested for retention of the same material several times. An experimental study of this sort yields very impure measures of retention after the first test, since all subsequent measures are contaminated by the practice the first test allows. As a practical matter, how-

ever, many studies of retention of school subjects have been carried out in this way.

A more adequate approach allows only one test of retention after learning. This can be accomplished by (1) repeated measurements on the same subjects, using different materials for each test, and (2) independent measurements of different subjects, using the same material. In the first method, the subjects may learn one set of nonsense syllables and be tested for retention after an hour, then learn another set of nonsense syllables of equivalent difficulty and be tested after two hours, etc. In the second method, the subjects are divided into several groups, each of which is tested after varying periods of time.

There are basic and profound differences between these two methods, as we shall see later in this chapter. Ebbinghaus and many subsequent investigators used the first method, but modern investigators prefer the second because it eliminates possible interaction between materials learned at different times.

If either the same subjects remembering different material or different subjects remembering the same material are to be compared, we must have some means of equating learning. This is generally done by using the same criterion of learning in all cases. Suppose that we want to find out how many nonsense syllables subjects can remember after 24 hours compared with the number they remember after 1 hour. We have two groups of subjects, both of which learn 15 nonsense syllables. We adopt a fixed criterion. Usually this means that we stop practice for each subject after he has correctly anticipated all 15 nonsense syllables twice in succession. The subjects may, of course, be given different amounts of practice before a test for retention, but the number of items *overtly* learned is held constant. Sometimes this is undesirable, and we adopt instead some criterion of number of trials. Thus we may give all subjects in the two groups, say, 10 trials of practice. In this case we hope that, on the average, the individual subjects in the two groups will have learned about the same number of items, though we actually allow the number of items overtly learned to vary freely.

In some cases we perform experiments in which there is no clear separation between learning and retention. Many experiments on such dynamic factors as motivation in learning and retention are of this sort. The simplest and most common example is the measurement of the "immediate memory span." The immediate memory span is simply the number (and kind) of separate items an individual can remember after only one exposure. Thus if we say the following series of numbers to someone, 4-7-2-6-3-8, and ask

RETENTION AND FORGETTING

him to repeat them, we are employing the procedure used in measuring the immediate memory span. We shall see the importance of this technique from examples given later in the chapter. Its limitation is that it does not allow a clear separation of learning and retention.

Measures of Retention

The measurement of retention is not so simple as it sounds. There are a number of different measures, and they do not always even approximate one another. We are justified in calling them all measures of retention in that they tend to be correlated, though some are much more sensitive than others, and on occasion, may reveal quite different things.

Direct measures. The simplest kind of retention measurement is a direct measure of performance. In many cases this simply means giving the subject one more learning trial. We may, for example, train a rat to run a maze by giving it 1 trial a day for 10 days. We then may "test for retention" by giving one more trial a month later. There may be no difference in the method of administering the trials; in the retention test, however, the emphasis is on the effects of the interval of time between the next to the last and the last trial. Thus we may compare retention after 30 days with retention after 1 day.

In verbal learning we may use an exactly analogous technique. We may ask subjects to learn a list of nonsense syllables by serial anticipation and test retention simply by presenting the list once more. In this test, the subjects must anticipate the next item. In verbal learning, however, we may also make use of a somewhat different method of directly testing retention, which minimizes the support the subjects receive from the context of what they have learned. This is the method of *recall*. After we have had subjects practice a list of nonsense syllables by serial anticipation, we may ask them simply to recall as many of the items in the list as they can. This eliminates much of the context provided by response-stimulus chaining in serial anticipation. The result is that such recall scores are generally lower than the anticipation or relearning scores.

Both recall and relearning scores may be expressed as a percentage of the number of items correct on the previous trial, which, if the subject achieves the ultimate criterion, will be all of the items. Thus if we ask a subject to learn a list of 20 syllables, and at recall he is able to recall only five of them, his recall, or retention score, is 25 per cent.

Savings. The direct methods suffer from a lack of sensitivity. It is possible, for example, that a subject may not be able to recall any nonsense

syllables at a test of retention, and his retention score is zero. Does this mean, then, that no effect of the previous learning remains? The answer is no, since a more sensitive method of probing may reveal some residual effects.

The method of *savings* is a much more sensitive method for quantitative tests of retention. It may reveal retention where none seems to exist by straight recall, and, indeed, it may reveal "negative" retention—it can show that an individual recalls less than if he had never practiced the material being tested.

As in the direct method, the subject is asked to practice an additional number of trials on what he has already practiced. The measure, however, is the percentage of time or trials *saved* compared with the original learning. It is given by the formula

$$\text{Per cent saved} = \frac{\text{No. trials to learn} - \text{No. trials to relearn}}{\text{No. trials to learn}} \times 100$$

Thus if it took 25 trials to memorize 15 nonsense syllables and 5 trials to relearn them after an interval of 1 hour, there remains a savings of 80 per cent.

Other methods. There are many variations of either the direct or the savings method. In addition, there is another important class of measures which can be summarized under the term *recognition methods*. These are very familiar, since one example is the multiple-choice test (invented by Ebbinghaus, by the way). This method has greater sensitivity than the direct measure, since it relies most fully on context. This is precisely its weakness for experimental work, however. It is furthermore dependent upon various artifacts of testing. If we simply ask someone whether or not he recognizes what we present to him, we run the risk of unreliable and false judgment. If we more strictly test an individual by asking him to pick out which of several items is correct, the probability of a correct selection on the part of the subject is at the mercy of the particular choices he has to make. We can vary the difficulty of an item simply by varying the incorrect alternatives. If we do this in a controlled manner, we may have a technique that is instructive. For most purposes, however, the method of recognition is limited to the applied problems of retention of school material and to the study of perceptual memory.

The study of perceptual memory sometimes introduces another technique, that of *reproduction*. This is a variant of the method of recall, in which an individual is asked to draw or otherwise reproduce something he has previously seen. As we shall see shortly, this method has likewise proved

to be very instructive. It is, however, much at the mercy of response factors (ability to draw, for example) and variability in the judgments necessary to scoring.

The Nature of Retention

We can now turn to studies of the course and nature of retention. In general, there are two things which investigators have been concerned about: (1) quantitative changes in the amount remembered with increasing time since learning, and (2) qualitative changes and reorganization of the material originally learned. We shall deal with the question of quantitative changes in memory first.

The Course of Retention

Let us look first at the studies of Ebbinghaus on the course of retention. He was much influenced by the notion that the elements of mental activity come in discrete packets called "ideas." The basic problem in the study of mental activity was how different ideas became associated and then disassociated. Ebbinghaus wanted to study the association of ideas in its simplest form, that in which completely new ideas are formed by the subject and the course of their retention examined by the experimenter. Hence the invention of the nonsense syllable.

We do not need to accept Ebbinghaus's devotion to the classical theory of the association of ideas, but one result of his views was that he studied retention of material in which, at the outset of the experiment, there was little or no associative value. In the light of our discussion about the effects of statistical dependency in material to be learned, we can see that his material was very difficult to learn and remember.

The results of Ebbinghaus's measurements of retention are shown in Figure 48. The most important feature of this curve is its extremely rapid fall. Most forgetting, or loss of retention, occurs in the first few hours after original learning.

The curve shown in Figure 48 has frequently been characterized as the typical retention curve, which in one sense it is. There are several important things to keep in mind about this curve, however. For one thing it was obtained with one highly practiced subject—Ebbinghaus himself. Another important thing to observe about this curve is that it is based on the savings measure, not simple recall.

In the older psychological literature, one may find many examples of the measurement, by various methods, of the amount retained over various pe-

riods of time and for various kinds of learning material. In general all these examples agree with Ebbinghaus's results in one respect—decline in retention is negatively accelerated (McGeoch, 1942; McGeoch and Irion, 1952). This seems to be true for the measurement of the retention of nonsense

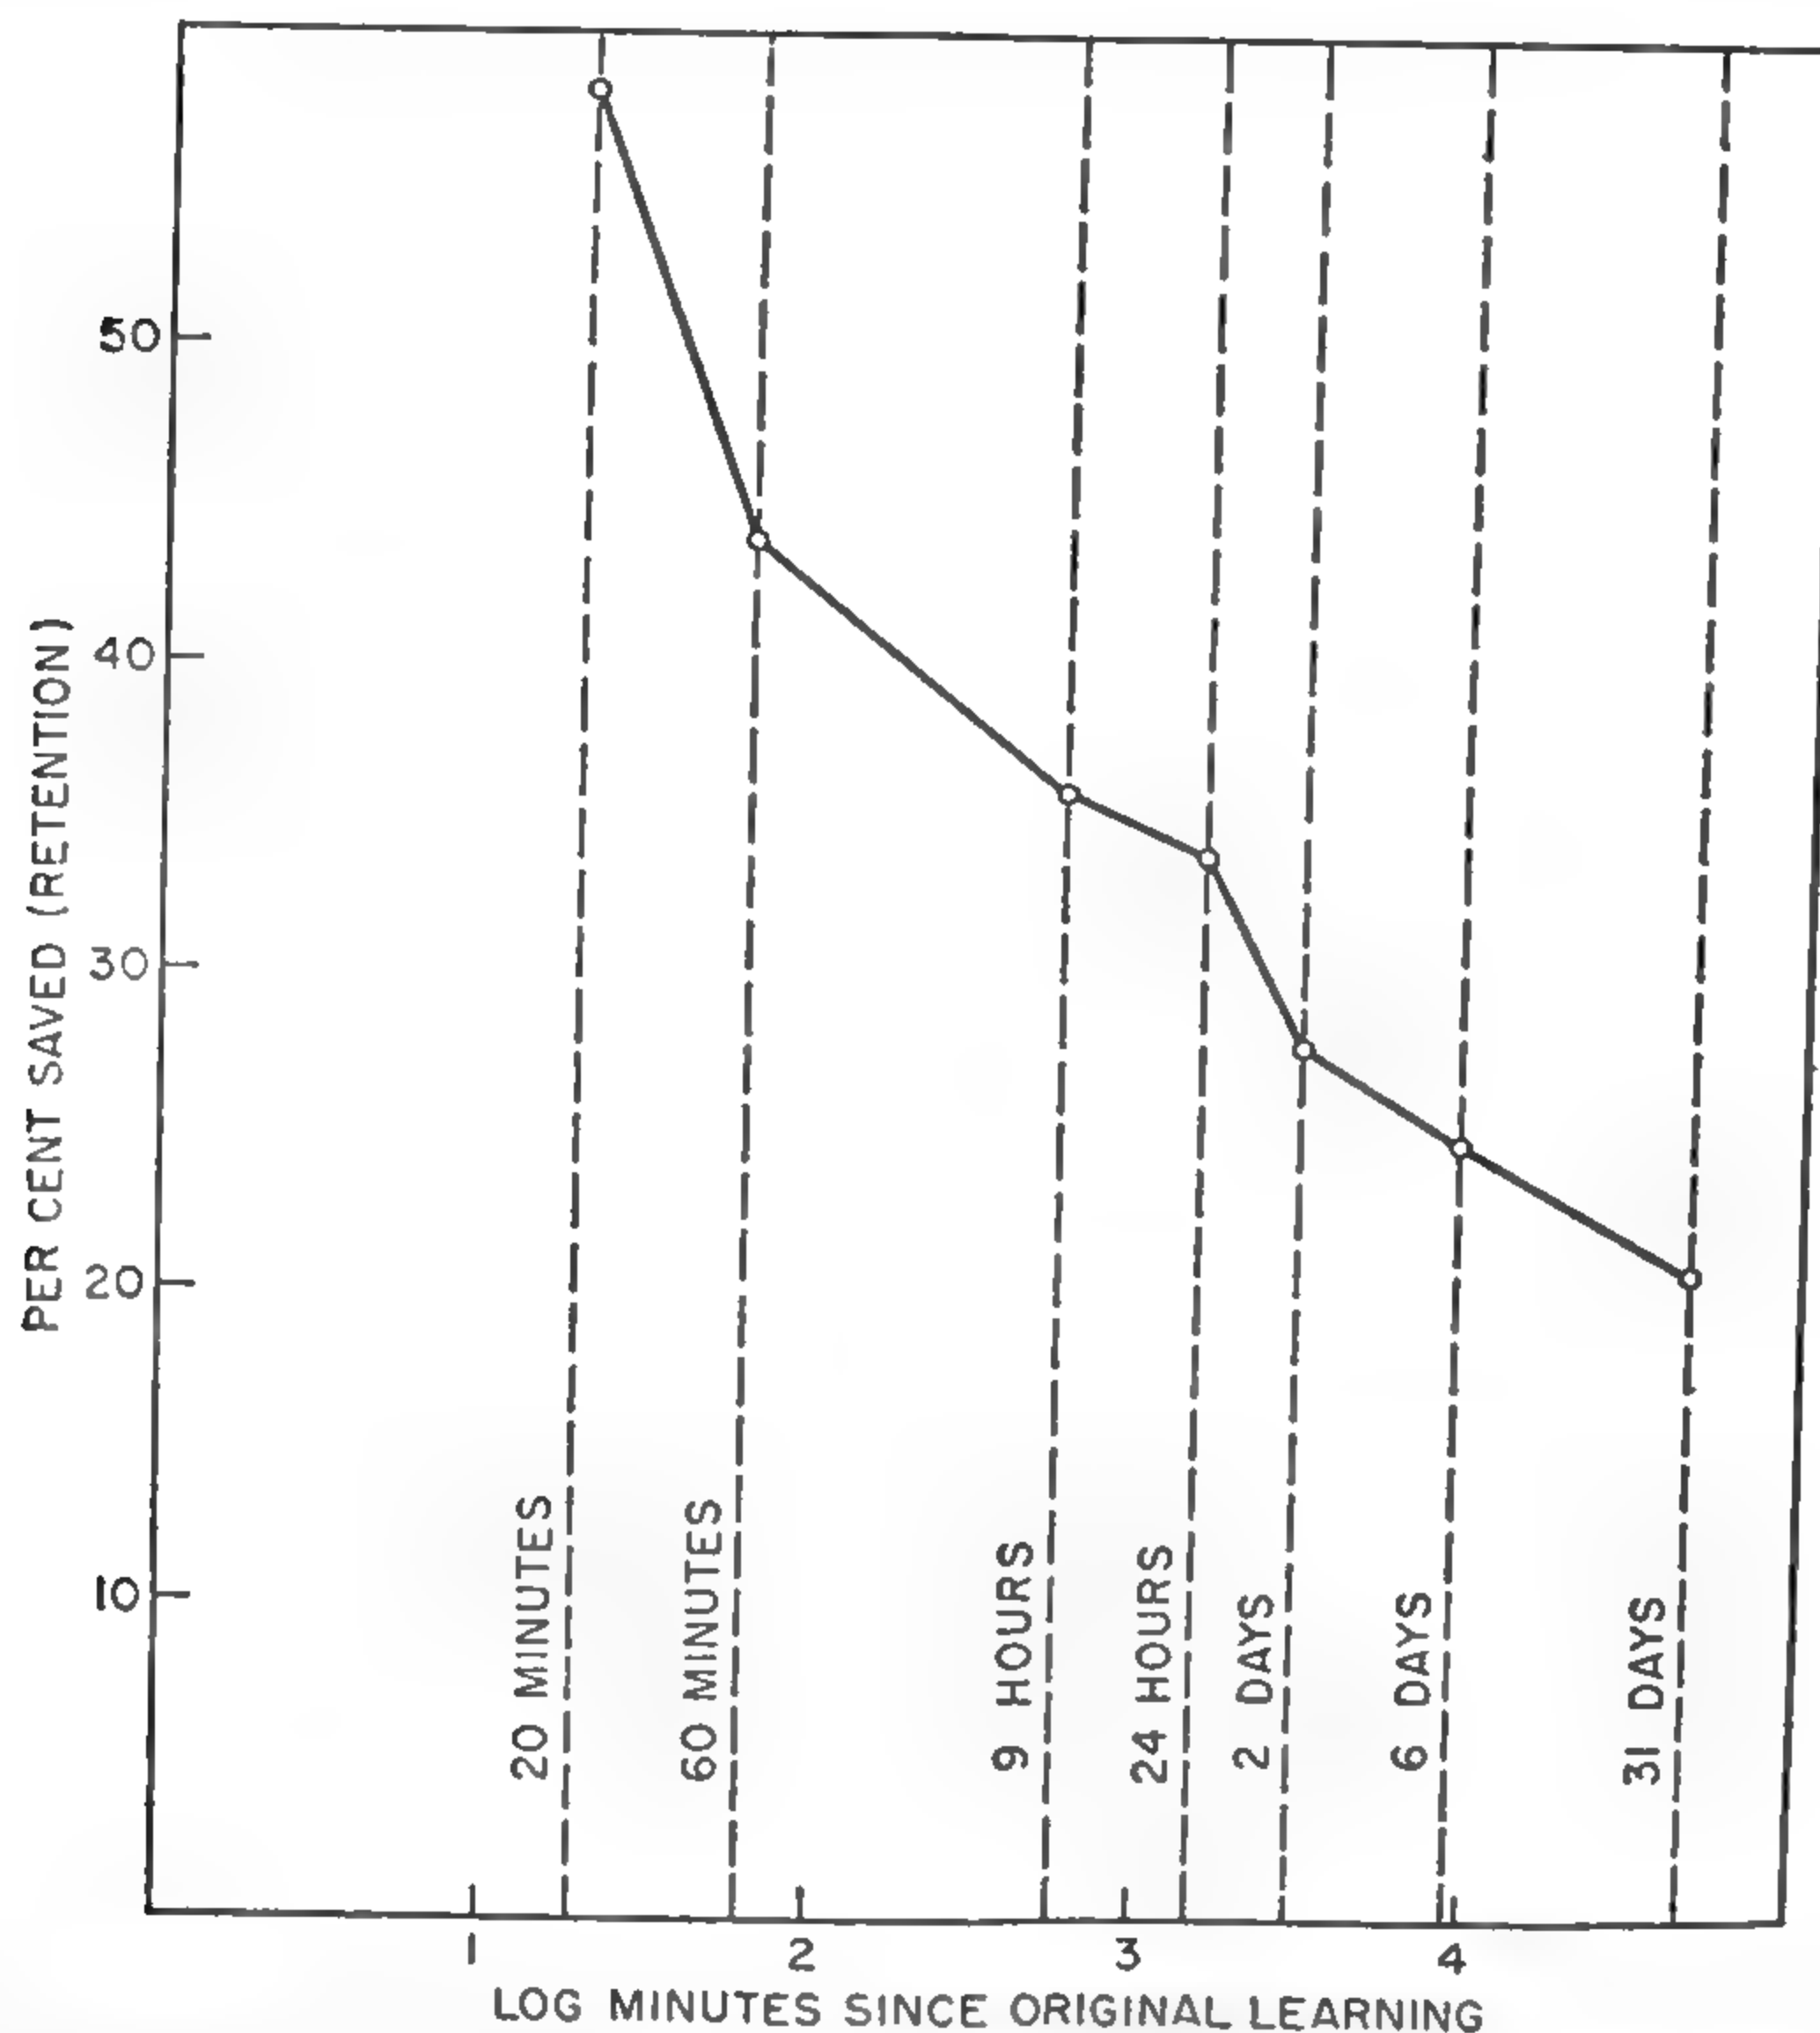


Figure 48. Retention (in per cent saved at relearning) as a function of time without practice. Notice that the curve is nearly a straight line plotted against log time. This means that most of forgetting took place in the first 24 hours. (Data from Ebbinghaus, 1885.)

syllables by other methods, the retention of meaningful material, and the retention of psychomotor skills. Of course, after special conditions of massed practice reminiscence may occur and lead to a temporary rise in the curve of retention. Under well-distributed practice, however, negative acceleration of the retention curve seems to be the general rule.

The rate and limit of forgetting is determined by many factors. The greater the degree of organization in material learned, the slower the rate of for-

getting and the higher the limit of retention. This in part is because of the greater redundancy of highly organized material. We know that when an individual is required to remember more material within a given context, there is lower retention measured by savings (Newman, 1939). Therefore, higher redundancy produces better retention principally because of the smaller number of things to be remembered. By the same token, memory for ideas in a prose passage is better than memory for factual detail (Briggs and Reed, 1943).

Another important determiner of the rate and extent of forgetting is amount of practice. It is well established that amount of retention is roughly proportional to amount of original practice (Ebbinghaus, 1885; Krueger, 1929). This relationship is sometimes expressed by saying that overlearning is beneficial to retention. Overlearning means continuing practice after some material has already been completely mastered. In general this point is well taken; practice beyond the stage at which an individual reaches a criterion of 100 per cent correct responses does improve retention, measured either by recall or savings. A more thorough study of this problem is needed, however. It is unlikely that practice beyond the criterion of one errorless trial would indefinitely continue to increase retention and that the relationship between practice and overlearning would be the same for all kinds of material. At present, however, we are not able to state clearly the limits of this relationship.

Retention and Speed of Learning

Most psychological textbooks, when discussing retention, attempt to "correct" a common misconception by pointing out that there is a positive relationship between speed of learning and amount retained. That is to say, contrary to folklore, which holds that someone who learns slowly retains well, the experimental literature apparently demonstrates that people who learn slowly retain comparatively little.

A paper by Underwood (1954b), which considers the problems of methods of measurement in comparing individual differences in speed of learning and retention, shows that the question is not so simple as it first appears. Underwood analyzed many data on rote serial learning, originally collected for other purposes. These, like earlier data, showed that there was nearly a zero correlation between time to learn and amount recalled; there was a high positive correlation, however, between time to learn and time to re-learn. Underwood further analyzed the problem by finding a number of items which were anticipated an equal number of times during learning by

fast and slow learners. On the average, of course, fast learners anticipated more frequently than slow learners, but it was possible to find some items in which the anticipations for fast and slow learners were equal. Even for these items, however, retention was better for fast learners.

This means, in effect, that any given correct anticipation strengthens an association more for a fast learner than for a slow learner. If this is so, degree of learning of the material is not equal before introduction of the retention interval. With this in mind, Underwood adjusted the association strength of items for fast and slow learners. When this was done, there was no difference in retention between the two groups.

The net effect of this analysis is to show that very likely the critical difference between fast and slow learners is simply the associative strength of particular items at the end of learning; fast learners have acquired more associative strength for particular items. It does not mean that fast learners have an inherent ability to remember things; it simply means that in a given amount of practice they generally acquire greater associative strength for the things they are learning. Perhaps they have more efficient means of coding (see below).

Reorganization in Retention

There is a good deal more to the problem of retention than simply whether or not responses are retained; they are reorganized in accordance with the individual's previous habits, motives, and other unique characteristics. Such reorganization is perhaps the most fundamental aspect of retention and we shall want to pay considerable attention to it, both in this and subsequent sections.

Clustering. One of the clearest and most important examples of reorganization of verbal material in memory occurs in the phenomenon of *clustering* discovered by Bousfield and his associates. Clustering can be demonstrated by the following experiment (Bousfield, 1953): Subjects listen to a list of 60 words. These fall into four categories—animals, names, professions, and vegetables. There are 15 words in each category. The 60 words are presented in random order, so that words from the same category occur together only by chance. Immediately afterwards the subjects write down, in order, all of the words they are able to recall. The two important results of this experiment were: (1) the subjects recalled more words than they would with lists of random words, and (2) they tended to recall words in the same category together.

As we have seen, in free recall of unorganized material subjects tend to

recall the last items of a list first and the remaining items in order of decreasing associative strength (Deese and Kaufman, 1957). When items can be categorized by the subjects, however, they reorganize the list so that they tend to give similar items together.

Bousfield points out that one of the principal determiners of whether or not a subject will recall a particular word in a list is the subject's previous experience with it. If it is a word that has occurred frequently in the subject's past it is more likely to be recalled than if it has occurred very infrequently.¹ This simple law of habit strength cannot account for clustering, however. Bousfield and Cohen (1953) have emphasized the importance of a second determiner of clustering. This they call a *relatedness increment*. They believe that words related to one another facilitate one another in recall by raising in strength a superordinate category of words to which all the words in a given category belong. The superordinate category is activated by a single high-strength word from the list; the category in turn then arouses the low-strength words in the list. For example, the occurrence of the word "leopard" activates the mediated response "animal," which in turn arouses the other subordinate members of this class (Bousfield and Cohen, 1955).

The important point about this notion is that the mechanism of clustering is not a direct matter of association. Rather it seems to be the result of an intervening mediational mechanism which is aroused by words on the list. Such a mechanism not only codes items presented to the subject in such a way that he can more easily recall them, but it reorganizes the order in which items are emitted in recall.

Number of categories in recall. An important problem in clustering concerns the number of categories into which the items on a test of recall may be classified. If one holds the length of the list constant and varies the number of categories, the number of items in a category decreases. What happens to recall under variation in the number of categories? One might expect that recall would decrease as the number of categories increased because too large a number of categories might exceed the immediate memory span (which is about seven items for adults). On the other hand, if the number of categories is reduced too far, there might be too many items in any one category for the immediate memory span, even taking into account

¹ The relationship between frequency of experience and habit or association strength is very likely negatively accelerated. That is to say, a word that is extremely familiar is not much stronger than one that is only moderately familiar, but a moderately familiar word is very much stronger than an unfamiliar word.

the advantages of coding. Perhaps one would expect some optimum number of categories for a list of a given length.

Actually, there has been little experimental work on this point. Mathews (1954) has studied frequency of recall for lists 24 items in length composed of two, three, and six categories, which in this case were composed of names of artists, athletes, musicians, etc. The experiment showed that recall increased slightly as the number of categories increased from two to six. This is perhaps what one would expect since six categories are not too great a burden on the immediate memory span (assuming that the subject could remember one name from each category). The names themselves, however, perhaps have low associative strength for any particular mediating superordinate, with the result that as the number of names gets too large and the number of categories gets too small, the probability of arousing a good superordinate is reduced.

The net result of all this is to suggest that we should be extremely cautious in generalizing the results on this point from any one limited experiment. If we knew more about the ways in which people formed their own mediating responses we could predict somewhat more adequately the specific effects of number and kinds of categories, but at present there seems to be no way of finding out about these.

The immediate memory span and recoding. We have mentioned immediate memory span frequently in the last few pages. It is of fundamental importance in the understanding of the organization of retention, and it deserves some attention in its own right.

The immediate memory span is the number of items of a particular kind which can be accurately reproduced after a single exposure. The immediate memory span for digits turns up as a standard item on nearly every intelligence test, and the norms for such tests tell us that the average adult is able to reproduce six or seven items accurately 50 per cent of the time. The immediate memory span is smaller with children and indeed, the fact that it regularly increases with age is one of the reasons it has a place in intelligence testing.

The size of the immediate memory span (in number of items correctly reproduced) varies with the material used for testing. For example, it is somewhat smaller for unrelated monosyllabic words than it is for decimal digits. Now this fact might lead one to conclude that the number of items of a particular kind that an individual can remember after just one hearing is determined by the amount of information per item in the testing material. Miller (1956b) has explored this hypothesis and found that it does not work.

We can think of the amount of information as the amount of uncertainty in any particular item about which we have to guess. Thus, if I ask someone to guess what word beginning with "A" I am thinking about, his uncertainty will be high, because there are so many words beginning with "A." If I ask him to guess what number between 1 and 10 I am thinking about, his uncertainty will still be high but not quite so high, since this time there are only 10 alternatives. If I ask him to guess which hand has the coin, the uncertainty is still less, since there are only two alternatives.

Information is an important concept in communication systems and computing devices, and when we say that a given communication system or computer has a given informational capacity, we mean roughly that it can reduce uncertainty about alternative answers to a problem only at a particular rate.

If our memory storage system has a limited capacity to take in information, then we should be able to remember more binary digits² than decimal digits and more decimal digits than words in a test of immediate memory. This is because there is *less* information in each binary digit than in each decimal digit and there is less information in each decimal digit than in each word.

We said a moment ago, however, that things are not that simple. Indeed, while our memory span for binary digits is in general larger than for decimal digits, and our memory span for decimal digits larger than for words, there is not nearly enough difference between them to make information recalled a constant. Indeed, we can remember only one or at the most two more decimal digits than monosyllabic words. Thus our ability to take in material and store it momentarily is not determined by the amount of information in the material. What then does determine the limits of immediate memory?

Miller (1956b) has made a start at answering this question. First of all, he says, we have to accept the notion that what people remember in the immediate memory span is *chunks* of things in which information has been encoded in various ways. For things like words, decimal digits, and indeed, any discrete series of items, there is not much difference in the number of items that can be recalled immediately. Therefore, it looks as though people remember a constant number of chunks or items *irrespective of the amount of information in each item*. Miller reminds us that if an individual has an immediate memory span of five words, we can guess that he has an immediate memory span of 15 or 20 phonemes, since each word, on the

² Binary digits are numbers in which only two digits are used, zero and one.

average, will have three or four phonemes. But, says Miller, people do not code by phonemes; they code by words, ideas, and sentences. Thus if people can remember an average of two sentences of simple prose, these sentences represent information coded into five, six, or seven chunks.

This process is called *recoding*. We hear something that has many chunks (say, words in a sentence) with relatively few pieces of information per chunk. We can take this sentence and recode it into fewer chunks (just the size to fit the immediate memory span) with more information per chunk. Thus we can get the "idea" out of a sentence. We do not know exactly what this recoding entails in every instance, but frequently it is probably a way of grouping the pieces of incoming information by applying some familiar name to the groups or pieces.

Miller reports an experimental attempt by one of his associates, Sidney Smith, to recode binary digits in an efficient manner. By regrouping the binary digits into octal units³ Smith could remember 40 binary numbers after a single hearing. How many items out of a sequence like this can you remember after one reading: 010110100100110101101011011100-1111010010? Perhaps as you read through the series, you could find a way of recoding these. Unpracticed people would probably try reorganizing by threes and read the sequence like this: 010-110-100-100-110, etc. This is not quite so efficient as recoding by assigning names to the possible groups of three. If you practice assigning names, you should be able to bring your immediate memory span for binary digits up to 36.

The point Miller wants to make is that memorizing of new material, the piling up, or storage, of additional verbal items in memory, is probably to a considerable extent simply the formation of chunks, or the reorganization of items into new groups. Finally, when we have reorganized the material into a small enough number of chunks we can recall all the material. Nonsense materials are harder to memorize than sensible verbal material simply because the latter comes already organized into chunks, whereas we ourselves have to form chunks for the nonsense material.

Clustering, the serial reorganization of material with statistical dependency, and related effects are probably all ways of reorganizing material into larger and larger chunks. These chunks most likely represent some mediational process, and we shall have more to say about this when we come to a discussion of problem solving and thinking.

³ Binary digits coded by assigning numbers 0 through 7 to groups of three. Thus 101 would be coded as "5" and 100 as "4."

Perceptual Memory

Let us now turn briefly to some problems in perceptual memory. Memory for visual forms, particularly abstract visual forms like circles, squares, and combinations of these, has been the special province of gestalt psychology.

The gestalt approach to perceptual memory. Gestalt psychologists have generally taken the view that there are "fields of organization" in perception and memory for perceptual events and that these are related to "physiological fields" in the cerebral cortex. They have considered the "memory trace" for a perceptual event to be autonomously altered by events intrinsic to the brain. In investigating the implications of these notions, they have studied the perception of figures in great detail and stated some rather simple rules that are supposed to govern the nature of changes in the memory traces for visual forms.

In general, gestalt psychologists have supposed that the memory traces for forms become better structured over a retention interval. Memories of forms are supposed to become simpler than the original perception. Experimental studies have not always confirmed the specific predictions various gestalt psychologists have made about the direction and extent of the changes in forms during a retention interval, and the actual events seem to be more complicated than gestalt notions allow.

There is one point which the gestalt psychologists have insisted upon, over the frequently vocal objections of more associationistic psychologists, and in this they have turned out to be absolutely right. This is the general proposition that there are systematic reorganizations going on during the retention interval. The memory trace for a perceptual event is not just a gradually fading picture of something we have seen in the past; it is something that is being actively reorganized and reconstructed.

It is likely that this reconstruction is much the same process as that we have just examined in the context of verbal retention. In other words, it is a reorganization that has as its roots certain recoding processes made possible by the fact that our perceptual world obeys probabilistic rules and that we have all had experience with these rules. Some of the most interesting studies on the way in which these recoding processes work have come from the classical investigations of Bartlett (1932), and we shall now look at some of his more important conclusions.

Bartlett's studies. First of all, Bartlett points out, when we are faced with an unfamiliar form, we do not just try to receive an impression of it

as it stands. We try to reconstruct it in such a way that it is easy to assimilate into our past experience with the world. There is, in Bartlett's words, an effort after meaning. This may be so simple as to be merely giving a name to the figure. Unless the figure is very simple and almost completely specified by the name, the figure will be schematized in memory. For example, the form in Figure 49 was frequently labeled by subjects in a perceptual-memory experiment as a "dog's head" (Deese, 1956). Such a designation will surely influence the way in which the form is retained.

Bartlett gives us several examples of the transformation of visual form in memory by the operation of schematic recoding. One of the best-known examples is reproduced in Figure 50. In this experiment the original drawing in the upper left-hand corner was presented to a subject. He was asked to reproduce this drawing from memory. The result is reproduction 1. This reproduction was then presented to another subject, and this subject was asked to reproduce the figure. We can clearly see how the form was transformed by the various schemata adopted by the different subjects. In general it is changed from the conventional ancient Egyptian symbol for "owl" to a picture of a sitting cat.

Bartlett describes at great length the various ways in which perception itself and remembering are transformed by schematization. The schematization partly depends upon the redundancy in visual forms (Attneave, 1954), and schematization based upon redundancy is likely to produce transformations in memory in the direction of simple, more symmetrical forms. Therefore it seems likely that some of the gestalt laws of perception and perceptual memory have as their basis the redundancy that exists in most forms of our



Figure 49. A "nonsense" form generally recoded by labeling when individuals are asked to remember the form. The most frequent name assigned this form by a group of subjects was "dog's head." (Deese, 1956.)

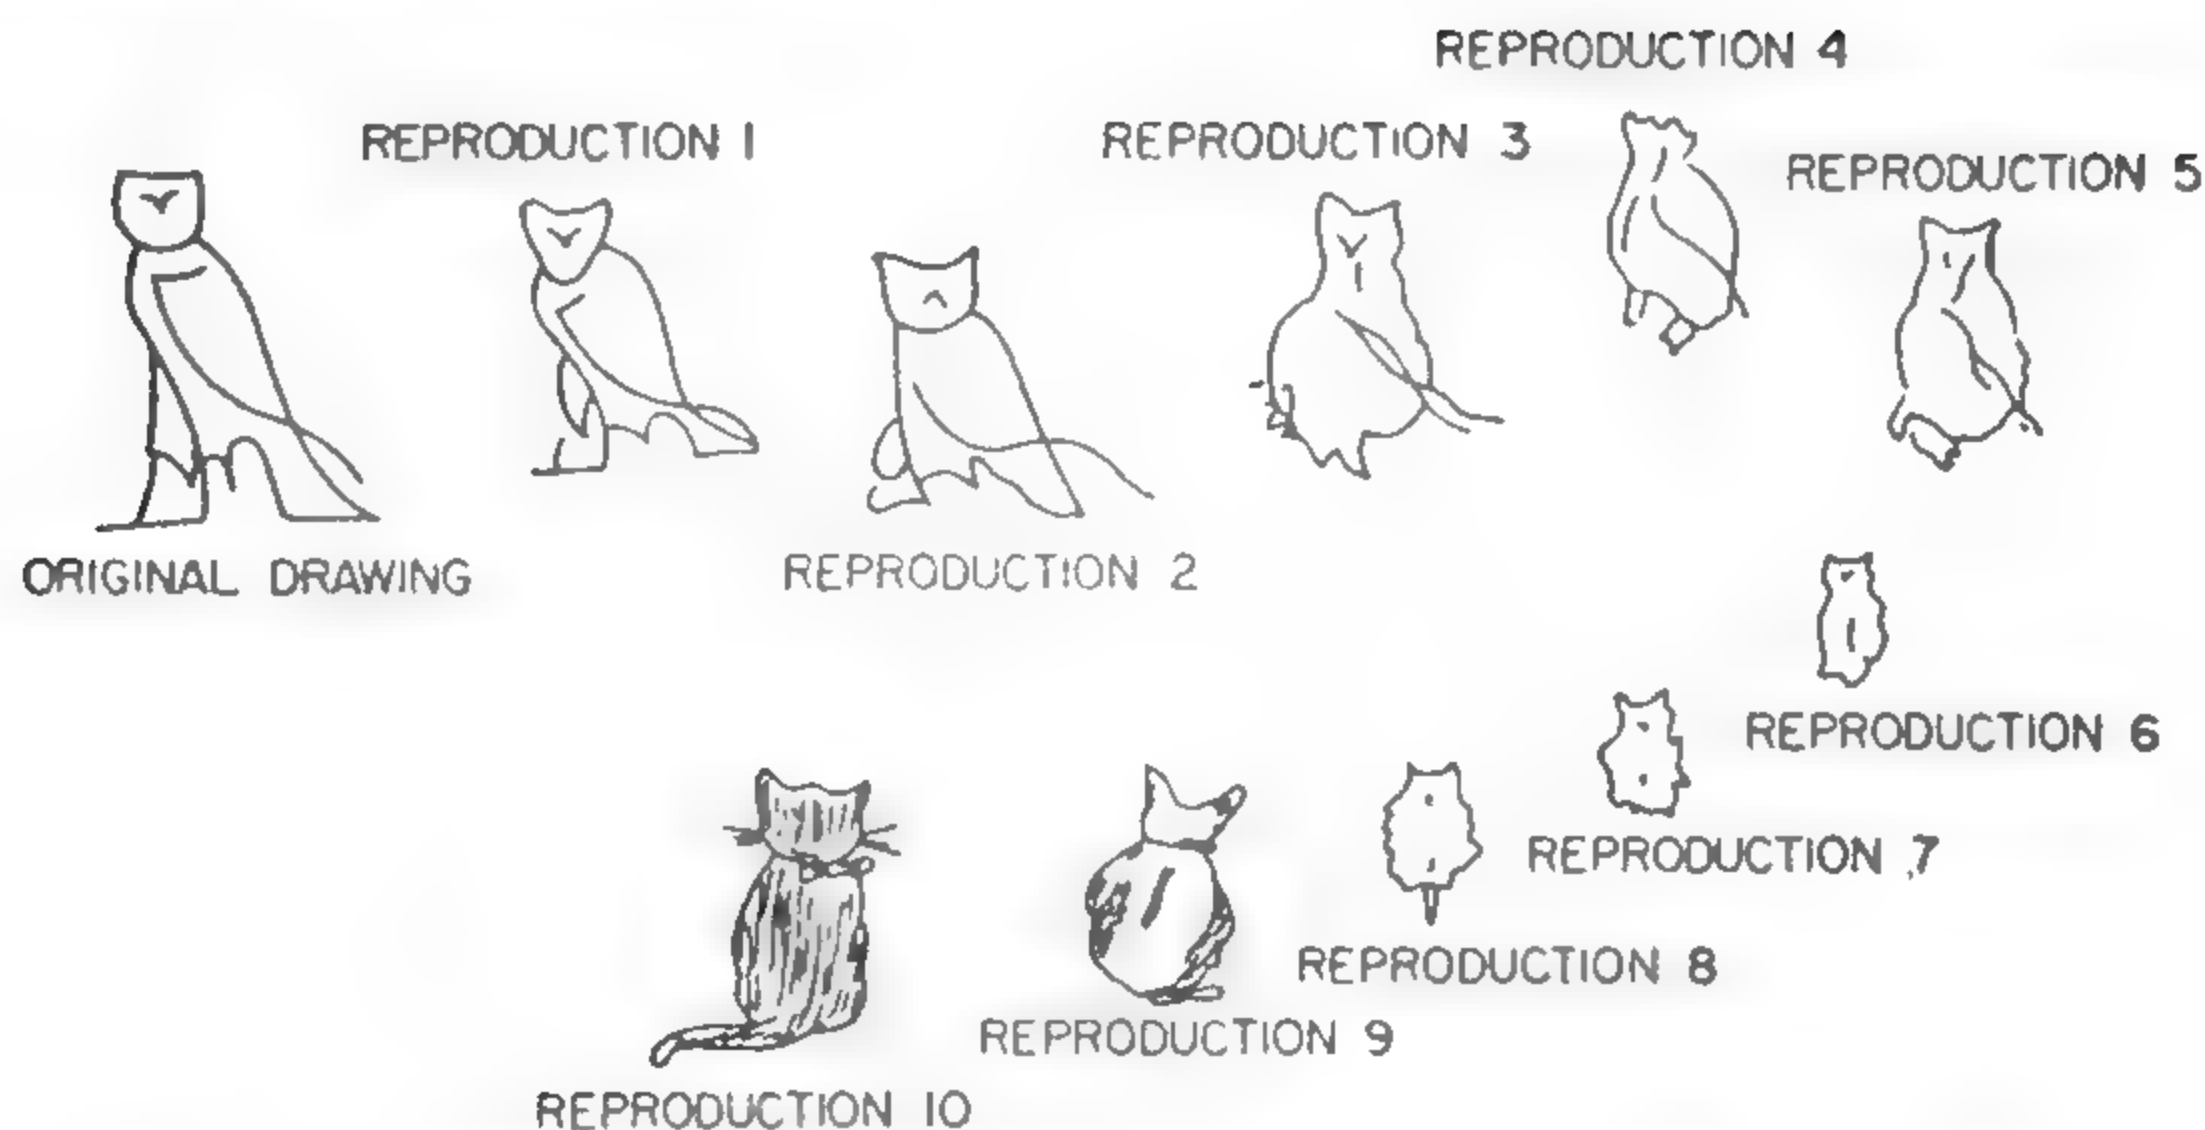


Figure 50. A form serially reproduced from memory by different subjects. The form is altered at each reproduction. The influence of labeling as a coding device is seen in the change from a conventional Egyptian symbol for an owl to a picture of a cat. (After Bartlett, 1932.)

experience.⁴ Redundancy, of course, means that some of the information presented in a visual form is superfluous; thus we can recognize a well-known person from a caricature of a few lines about as well as we can from a photograph.

Ways of schematizing. We have already mentioned some of the ways in which forms may be schematized for visual memory. They may have names attached to them or be reconstructed by some simple principle of design, particularly if there is high redundancy in them. Actually the ways in which schematizing may occur are numerous, and it will only be worthwhile to illustrate a few.

Simple forms may be reproduced even if they cannot be well schematized, much as a list of random digits may be reproduced in the immediate memory span. If nonsense forms become too complicated, however, they must be learned by rote, and the length of time this takes is about the same as for any other kind of nonsense material with an equal number of items (Piéron, 1920). The determinants of memory for any set of perceptual events, however, are quite complicated.

For example, if subjects are required to recognize rather than reproduce forms, recoding is likely to be quite different. They may recode by naming, but if names are difficult to assign to the forms, they may try to identify them by some isolated segment, characteristic of a form. If recoding is done this way it may turn out that a complicated form will be easier to recognize than a simple one, because the former will have a greater probability of

⁴ See, for example, Attneave (1955) on symmetry.

having some really distinctive, isolated segment (Deese, 1956). On the other hand, if the subjects are required to view too many different forms, recognition for simple forms may be better than for complicated forms because the number of distinctive segments the subject may have to try to retain for the complicated figures exceeds the immediate memory span.

The importance of the limitation of recoding is illustrated in an experiment by French (1954). This investigator had subjects identify patterns of dots from memory. The subjects were required to learn a name associated with each pattern of dots, which were arranged randomly so that recoding by familiar names could occur only by chance. The results were that subjects could more easily recognize patterns with six to eight dots than they could patterns with fewer or more dots. Ease of identification was associated with patterns of dots arranged either symmetrically or in linear arrays. In other words there was better perceptual organization in the patterns of six to eight dots. With fewer dots there was less likelihood of organization (and hence, of recoding), and with more dots the patterns became too cluttered.

The ability to recode our visual experience, for one thing, appears to be a function of the frequency of exposure to certain forms. In one experiment (Arnoult, 1956) subjects were exposed to nonsense forms a varying number of times. These same subjects were later asked to rate these forms for familiarity. There was positive relationship between the frequency of exposure and familiarity, much as there is with unfamiliar words (Noble, 1954). Furthermore there were positive relationships between the recognizability of the forms and their frequency of presentation and rated familiarity.

Dynamic Factors in Memory

If one reads enough books in psychology he will get used to seeing the phrase *dynamic factors*. Generally this refers to individual differences in personality, motivation, emotionality, and related factors. Nearly all psychologists assume that such factors influence memory. Sometimes particular theories are quite specific about the ways in which dynamic mechanisms may influence memory, as is the case with the Freudian notion of repression. More often, however, there is simply general agreement that the specific personality mechanisms of an individual will influence things that he remembers.

Individual styles in recoding. Again, we return to Bartlett's studies (1932). One of the most influential aspects of these was his observation

that the nature and direction of schematizing, or recoding, is determined by the personal characteristics of the individual doing the recalling. The extent to which personal characteristics determine the nature of schematizing, of course, depends upon the material being recalled. In one of Bartlett's studies, he asked subjects to indicate their memories for pictures of faces by describing the faces after the subjects had seen them. This task appeared to be particularly susceptible to personal attitudes and motivation. A particular face would be described in terms of the conventional attitudes toward the type it represented. Thus, a picture of an enlisted man in the army would be described as representing a young, good-humored person, etc., and a picture of an officer would be described as stern, elderly, etc. Obviously, the attitude aroused by the pictorial schema or the name assigned to the picture by a particular subject influenced his subsequent description of the picture.

We can generalize this kind of observation by saying that the memory for a particular event is influenced by the schema the person uses to carry it in memory. The schema, in turn, is at the mercy of the individual's motivations and attitudes, some of which are specifically aroused by the original event to be remembered and some by a label or mediating reaction given to the event to be remembered.

Such observations have been extremely important for certain problems in social psychology, and indeed, Bartlett's studies have been as influential in this area as they have been in the psychology of learning. They provide an important point of contact between the principles of social and of individual behavior.

Selective forgetting. Retention—or the loss of it, forgetting—seems to be selective in nature. That is to say, we remember some things and forget others, not because of the inherent associative strength of the material learned, but because they are related to our motivations and attitudes. In brief, it has commonly been assumed that people remember those things in which they are interested.

One of the best-known studies usually cited to support this general proposition is by Levine and Murphy (1943). These investigators presented subjects with material dealing with the Soviet Union that was either favorable or unfavorable. Before the study, the subjects had been divided into two groups on the basis of their attitudes toward communism—procommunist or anticomunist. Levine and Murphy obtained learning curves and retention curves for these two groups. It turned out that there was a marked relationship between the attitudes of the subjects and amount learned. Thus,

the procommunists were able to reproduce, during learning, a larger percentage of the separate ideas from the procommunitistic material than from the anticommutistic material. Unfortunately, however, because of the difference in learning, the retention curves started at different places, and hence it is extremely difficult to tell whether there was selective forgetting or not. In other words, there was selective learning, but this tended to mask the possibilities of finding selective retention.

Since selective learning seems to happen so much of the time, it is extremely difficult to tell whether or not selective forgetting also occurs. One experiment at least suggests that it does. In this experiment (Taft, 1954) Negro and white boys were asked to recall material read to them aloud; the material contained favorable and unfavorable, ambiguous and neutral items about Negroes. On an immediate recall test the white boys recalled about as many neutral items as did the Negro boys, but the Negro boys recalled more emotionally toned items, both favorable and unfavorable, than did the whites. The greater recall for the emotionally toned material on the part of the Negro boys was probably influenced in part by the fact that the experimenter was white and was an authority figure of considerable importance.

Thus far this experiment tells us little that is new. However, the experimenter went one step further and asked for a delayed recall three days after the original recall. On the delayed recall the Negro youths were still superior to the whites in the recall of favorable material, but they showed considerably more forgetting of the unfavorable items than did the whites, so much so, as a matter of fact, that the whites were slightly superior on the recall of unfavorable items.

Why did the Negro youths forget the unfavorable items? Very likely because they fitted less well into their emotionally determined schemata about successful Negro baseball players (this is what the material was about). In addition, however, there may have been deliberate suppression of the material (unwillingness to tell the experimenter about the unfavorable material) or repression. This last mechanism is extremely important, and we shall deal with it later.

Once again we must emphasize the importance of the present assimilation of what the individual perceives with what he has learned in the past. This will determine what an individual learns about a new situation and probably what he retains of it. Aside from the emotional and attitudinal content of this problem it is not basically different from the general question of transfer of training. A new situation arouses a set of well-established

and preferred associations in the individual, and it is to these responses that parts of the new situation become selectively attached. When the time for recall comes, the individual is able to bring out those parts of the new situation that can readily be aroused by strong associations—those things that have become assimilated into the psychological life history of the individual.

THE MECHANISMS OF FORGETTING

We have seen that individuals tend to remember things which are related in a certain way to what they have learned in the past. In most of the experimental literature we have described in support of this general proposition, the amount of learning was uncontrolled. Thus we have been dealing with both learning and retention. If we have established, however, that some responses have all been equally well learned, what determines how many, and which, of these responses will be forgotten? The answer to this question is the subject of this section. We shall deal with studies in which investigators have deliberately tried to produce forgetting and with the theories that have arisen out of these experiments.

An assumption frequently made is that forgetting is simply a matter of lack of practice. A learned act, it is argued, disintegrates because time passes without any practice, much as a footpath will be overgrown with weeds if no one walks along it. This idea is reasonable enough, for in the natural course of events, our learning seems to slip away from us with the passing years. The schoolteacher who returns to the university for a summer session after five years of teaching the sixth grade makes a determined effort to impress her French instructor with the fact that she has not looked at French in years and consequently will not remember much.

Arguments against the principle of disuse. Forgetting believed to be a result of lack of practice usually is called *disuse*. There are many grounds for dismissing disuse as a serious notion in the psychology of learning. McGeoch (1932) pointed out some of these in a classical paper on the nature of forgetting. The most important point that he made is that disuse does not explain forgetting if disuse only implies the passage of time, for time in itself does not cause anything. Events happen in time; certain conditions change over time, and it is these that provide explanations. Thus we find that forgetting is not determined by passage of time alone, but by the nature of the events which fill a time interval. An analysis of these events will lead to an understanding of the nature of the forgetting process itself.

Transfer and Forgetting

If we once admit the proposition that forgetting is determined by events filling a time interval, it immediately becomes apparent that the experimental designs for the study of transfer of training are also those for the production of forgetting. Negative transfer, then, is admitted to be one of the mechanisms of forgetting.

Retroactive Inhibition and Proactive Inhibition

Experimental designs. As we saw in the chapter on transfer there are two types of basic experimental designs for the study of transfer. The retroaction experiment measures the effect of an interpolated activity on the retention of some previously learned act; the proaction experiment measures the effect of an older learned act upon the learning and retention of a new act. The basic designs are as follows:

Proaction

Experimental group

Learns task 1learns task 2 (recalls task 2)

Control group

Restslearns task 2 (recalls task 2)

The influence of task 1 upon task 2 can be seen by examining either the learning of task 2 or the retention of that task sometime after its original learning.

Retroaction

Experimental group

Learns task 1learns task 2 tested on task 1

Control group

Learns task 1rests tested on task 1

The influence of task 2 is measured by a retention test (usually recall or relearning) of task 1.

If the experimental condition impairs the retention of an old task (retroaction) or the learning or retention of a new task (proaction), negative transfer or inhibition has taken place. Hence the special names retroactive inhibition and proactive inhibition for the negative side of transfer. Most investigators have considered retroactive inhibition to be the more

important condition for the mechanisms of transfer in forgetting, but more recently proactive inhibition has been assigned an important role also.

In order to see how the conditions of negative transfer have been applied to a theoretical analysis of forgetting, it will be necessary for us to examine some of the principal experimental findings in negative transfer. In the next few pages we shall look at some of the more basic experimental variables.

Negative transfer and stimulus-response relationships. In the chapter on transfer we examined a diagram, the transfer surface, which implied that the major conditions for negative transfer between two tasks were (1) common stimulus elements between the tasks and (2) different response elements. The diagram implied that, holding stimulus similarity between the tasks constant, transfer could be varied from positive to negative by changing the responses in the two tasks from being identical or very similar, to being very different from each other.

This diagram was drawn by Osgood (1949) to summarize the findings of many different experiments. In other words, the basic relationships presented in the diagram are based upon empirical research. A large number of experiments, using both the proaction and the retroaction design, have shown that when stimuli are the same, but responses different, negative transfer results (Bruce, 1933; Gibson, 1940; Bugelski, 1942; Underwood, 1945; Osgood, 1946; Morgan and Underwood, 1950).

Thus, a safe generalization about paired-associate learning is that when the same stimulus items and different response items are used in two tasks there will be negative transfer.

The transfer surface also implies that the direction and amount of transfer depend upon the *degree* of response similarity. As this similarity goes from a maximum (identity) to a minimum (as Osgood suggested, opposite words), transfer goes from maximum positive to maximum negative transfer. The empirical evidence, however, does not entirely support this relationship between response similarity and transfer.

Some experiments appear to uphold the notion that transfer varies continuously as a function of response similarity (Morgan and Underwood, 1950). Other experiments cast doubt on the generality of this relationship. For example, Bugelski and Cadwallader (1956) performed an experiment designed to test all of the implications of the transfer surface. In this experiment the greatest amount of retroactive inhibition was for responses in recall which were somewhat similar to responses in interpolated learning, rather than for responses which were neutral or opposite.

This is a complicated and difficult issue to settle. It may be partially the result of the amount of interpolated learning, as it was in an experiment (Deese and Hardman, 1954) cited in the chapter on transfer. On the other hand, it may be the result of the way in which the subjects are able to re-code responses. For example, if during interpolated learning, the subjects become aware that the responses required are all antonyms of the responses required in original learning, the net effect, both on interpolated learning and recall of the original list, may be positive transfer. At recall the subject remembers that one of the words in the interpolated list was "swift" and quickly thinks of the antonym "slow." If, however, the words on the two lists are unrelated to each other, no simple recoding will do. Perhaps maximum negative transfer would be obtained in situations in which responses are simply unrelated or are so similar that the subject cannot remember which word belongs to which list. Suppose, for example, that a subject remembers that "swift" was a response item for one task and "fast," for the other; since they are similar in meaning, however, he has difficulty in recoding them to their proper lists. He is likely, therefore, simply to guess one or the other, and there is a 50 per cent chance of making a mistake.

This matter of the relationship between recodings available to the subject and response relationships in negative transfer still needs a thorough experimental exploration.

A comparison of proaction and retroaction. Several investigators have asked the question, Is retroactive inhibition greater than the proactive inhibition of retention when the same materials are used to test both? This is important, because we wish to know whether these two situations are measures of the same psychological effect or whether one involves something more than the other.

One classical experiment on this comparison (Melton and von Lackum, 1941) used nonsense syllables and the technique of serial anticipation. Under these conditions retroactive inhibition was greater than proactive inhibition of retention. That is to say, an interpolated task has more inhibiting effect on the retention of an original task than an original task has on the *retention* of a task learned second. This finding will not violate the layman's prejudices, since in the proactive experiment the more recent task is measured for retention, and in the retroactive experiment the earlier task is the one measured.

If, however, a time interval is allowed between original learning and the test for retention, the difference between proaction and retroaction disappears (Underwood, 1948). This again, perhaps, will not surprise the lay-

man, since the order difference between the tasks in proaction and retroaction will be slight compared with the time difference of a long retention period—if there is a long enough time lapse, the order of the tasks becomes unimportant. Thus, while there is a difference between retroactive and proactive inhibition when the retention test is immediate or follows soon after original learning, the difference disappears if the retention test is delayed. This effect is largely the result of *recovery* from retroactive inhibition rather than increase in proactive inhibition (Underwood, 1948), though some increase in proactive inhibition probably does occur after a long retention interval (Duncan and Underwood, 1953).

Proaction, retroaction, and degree of learning. If we compare the effects of learning and retention of a second task after differing amounts of practice on the first task, we find that both positive transfer and proactive inhibition of retention increase as a function of the degree of prior learning (Atwater, 1953). That is to say, the degree to which the acquisition of a new task is aided by previous learning depends upon the amount of practice on the original task (such aid is probably largely “learning how to learn”), but by the same token, the greater the learning of a prior task, the more it will interfere with retention of the second task.

It is not clear, however, in what way amount of practice on the preliminary task and on the second task interact. We do know that “overlearning” of a single task makes it more resistant to ordinary forgetting. If the preliminary task is “overlearned,” the responses in it may be so thoroughly associated with one another that the learning of a new set of responses will not interfere with them.

In the retroactive design increased practice on an interpolated task has a more complicated effect, at least in rote verbal learning. With a fixed number of trials on original learning (nonsense syllables), an increase in practice on the interpolated task first results in an increase in retroactive inhibition. After the interpolated task has been practiced to a much higher level than the original task, however, further practice on the interpolated task seems to decrease the amount of retroactive inhibition (Melton and Irwin, 1940; Thune and Underwood, 1943). Thus, it is possible that if the tasks are highly differentiated by overlearning either the original or the interpolated task, retroactive inhibition may be less, at least in an immediate test of retention.

There is still some doubt, however, about how general this finding may be. For motor learning, at least one experiment (Lewis, Smith, and McAllister, 1952) shows that the amount of retroactive interference is simply

inversely related to criterion level on an interpolated task. Another experiment on rote verbal learning (Briggs, 1957) shows that retroactive inhibition increased to an asymptotic limit with an increase in practice of the interpolated task, but did not decrease. This, however, may have been simply because practice on the interpolated task was not carried far enough.

In general, we can conclude that inhibition (either proactive or retroactive) grows with an increasing amount of practice on the interfering task, but it is likely that there is a limit to this increase; it may even be that with very high amounts of practice on the interfering task, the inhibiting effects may decline slightly.

Theories of Retroactive and Proactive Inhibition

What factors are responsible for proactive and retroactive inhibition effects? If we regard the experiments on these effects as models of the experimental production of forgetting, then a reply to this question would provide us with some of the factors responsible for the forgetting process. The most frequently given answer at the present time is that competition between responses is one reason for both proactive and retroactive inhibition. Before we examine some of the implications of this notion, however, we shall examine one of the other possible theories of inhibition, one which has been conceived of as interruption, or disruption, of a memory trace.

Interruption of a memory trace. The classical theory of retroactive inhibition is usually known as the "perseveration theory," and the essence of this is that after a series of responses has been practiced, the effects persevere for a period of time (presumably in the central nervous system) and that any other activity will interfere with these effects. This notion, incidentally, has been advanced to account for the distribution of practice effect.

As far as retroactive inhibition is concerned this idea is of limited use. For one thing, it implies that retroactive inhibition should occur most strongly when a test for retention of interpolated learning is given immediately after interpolated learning, and indeed, it should not occur at all if the test is delayed too long. This is contrary to fact (Postman and Alper, 1946).

Also against this notion as a general explanation of retroactive inhibition is the fact that it does not predict the rather important effects of similarity (both with respect to stimuli and responses) upon retroactive inhibition. Furthermore, it does not allow for an explanation of retroactive inhibition and retroactive facilitation by the same mechanisms. Finally, it must lead one to the conclusion that retroactive and proactive inhibition of retention

are essentially different phenomena, since the perseveration notion cannot explain the latter.

These facts have led contemporary theorists to abandon perseveration as a general theory for retroactive inhibition. Nevertheless, there may be some facts about learning which demand a construct such as perseveration, or at any rate, suggest the role of interruption of the memory trace immediately after it is formed as a factor in forgetting. For one thing there is the existence of the immediate memory span. Why is our memory span for unrelated items limited to about seven items? Why cannot we hear a long string of numbers and then be able to recite them as well as we can when we split the string in two and recite the two halves separately? The fact that we cannot suggests that the later items we hear somehow interfere with our memory for the first items before we get a chance to emit or recode them.

Hebb (1949) has suggested that there are two phases to the neural trace of something learned, dynamic and permanent. The dynamic phase lasts but a short time, and, Hebb argues, it depends upon something like active recurrent circuits in the nervous system. If anything acts to interfere with the event to be remembered before it is transformed into a more or less static memory trace, it is disrupted.

Such an effect, however, probably has little to do with retroactive inhibition as it ordinarily occurs. The "traces" established in a retroactive inhibition experiment are durable and exist much longer than the few seconds in immediate memory span. Therefore the notions about consolidation or perseveration of traces are not useful as explanatory constructs for the retroactive and proactive inhibition effects. They may be useful in explaining some of the fundamental facts of *learning*, but thus far there has been little interest in theories of learning built on the notion of consolidation of traces.

Competition theory. The most specific of the several related concepts that have been advanced to account for retroactive and proactive inhibition is the hypothesis of response competition. This asserts that there will be mutual interference between two sets of responses to the extent that the responses from the two sets compete with one another at recall. The extent to which one set of responses will compete with another depends upon the relative associative strength of the sets of responses to common stimuli and the inability of the individual to recode the responses in such a way as to give them distinctive contexts.

The effects of competition may be expressed in several ways: (1) The individual may be able to recall both sets of responses but may not be able to recognize which one is correct. (2) The individual may be able to recall

only one of the alternative responses and not the other. Competition here may be indicated either by (a) the recall of the incorrect response or (b) a longer reaction time for the recall of the correct response. (3) Finally, competition may depress all alternative responses, so that for everything except probing by a forced-choice method, the individual may not be able to recall either set of responses. Indeed both sets may be depressed enough for a third "irrelevant" set of responses to occur.

Competition can occur only in a situation in which one of the alternative sets of responses is correct. If either one of the two alternative sets will do in response to a particular situation, competition, by definition, does not occur. Furthermore, competition will not occur when the individual can recall both sets of responses and both are acceptable. Neither of these conditions usually arises in formal tests of retention. For example, in answer to a completion item such as "the receptors for color vision are the _____," an instructor would accept only "cones," not "rods or cones."

Implications of the competition hypothesis. The competition hypothesis, even in the bare outline presented above, has several rather specific implications. For one thing it implies that the degree of competition between habits will depend primarily upon the associative strength of the habits for identical or related stimuli. In other words, the stronger a competing response, the more it will compete. Furthermore, if associative strength between the common stimuli and the two sets of responses is equal, then temporal order should have no effect. This means, all other things being the same, that there should be no difference between retroactive and proactive inhibition of retention. Likewise, there should be no difference in amount of retroactive inhibition produced by variation in the temporal position of interpolated learning.

Another set of implications of the competition hypothesis is that both proactive and retroactive inhibition should be determined by the similarity relationships between tasks. For example, if subjects learn two distinct sets of responses to stimuli, competition at recall will be greater if the stimuli with which the responses have been associated are identical rather than similar. In other words, something like stimulus generalization will operate.

The competition hypothesis does not have much to say about response similarity, however, except that if responses are so imperceptibly different that either response will be accepted as correct, then positive rather than negative transfer will occur. Once the responses are distinct from one another, however, the role of degree of response similarity in the competition hypothesis is not clear, since recoding is possible. Perhaps this is just

as well, for we have seen that the experimental evidence on negative transfer and response similarity is not unequivocal.

On a number of these points the data are not really in disagreement with the implications of the competition hypothesis. To be sure, in a test of immediate recall there is likely to be a difference between the retroactive and the proactive inhibition of retention, but this disappears in a delayed recall test. The difference on initial recall may have to do with subtle factors of set and warming up. Certainly, in the main, the transfer surface is in accord with the implications of the competition hypothesis. The data on the temporal point of interpolation in retroactive inhibition are contradictory and confused, which is perhaps exactly what one would expect from the competition hypothesis.

Since competition is determined by the associative strength of the competing activity, we should expect that retention in retroactive inhibition would vary inversely with the amount of interpolated learning and that retention in proactive inhibition would vary inversely with the amount of learning on the original task. Again, this appears to be the case, with the possible exception of high degrees of interpolated learning in retroactive inhibition. The slight drop in retroactive inhibition that occurs in such a situation may be the result of differentiation of the two response systems through efficient recoding.

Direct transfer and the competition hypothesis. Perhaps more important than these indirect implications of the competition hypothesis are direct indications of overt negative transfer. Direct expression of competition can occur in two ways: (1) the reaction time for correct responses may be increased by the need to suppress incorrect alternatives at recall, and (2) overt intrusions of responses from the competing set of responses may actually displace correct responses in recall.

Both these effects occur. Postman and Kaplan (1947) have observed that reaction time during relearning of the original responses is increased as the result of interpolated learning. From this increase it is easy to infer that the correct responses are inhibited for a time by an occasional tendency to say the appropriate word from interpolated learning. Everyone has had the experience of not quite being able to recall something without collecting his thoughts for a moment, and this generally means picking the correct response from a number of alternatives.

Of critical importance to the competition theory has been the observation of direct intrusion of responses from the competing set of habits. Ordinarily, in a rote learning experiment, most of the "errors" are simply failures of

the subject to respond within the assigned time limit. Sometimes, however, subjects will respond with inappropriate responses, and these have provided a direct index of competition. The very fact that such intrusions occur is direct and conclusive evidence that competition is responsible for at least part of interference with recall.

Melton and Irwin (1940) noticed that the frequency of intrusions of interpolated responses during the recall of an original task was greatest when the amount of practice on the original and interpolated tasks was about the same. If there was more practice on the interpolated task, interpolated responses tended to intrude less frequently during recall of the

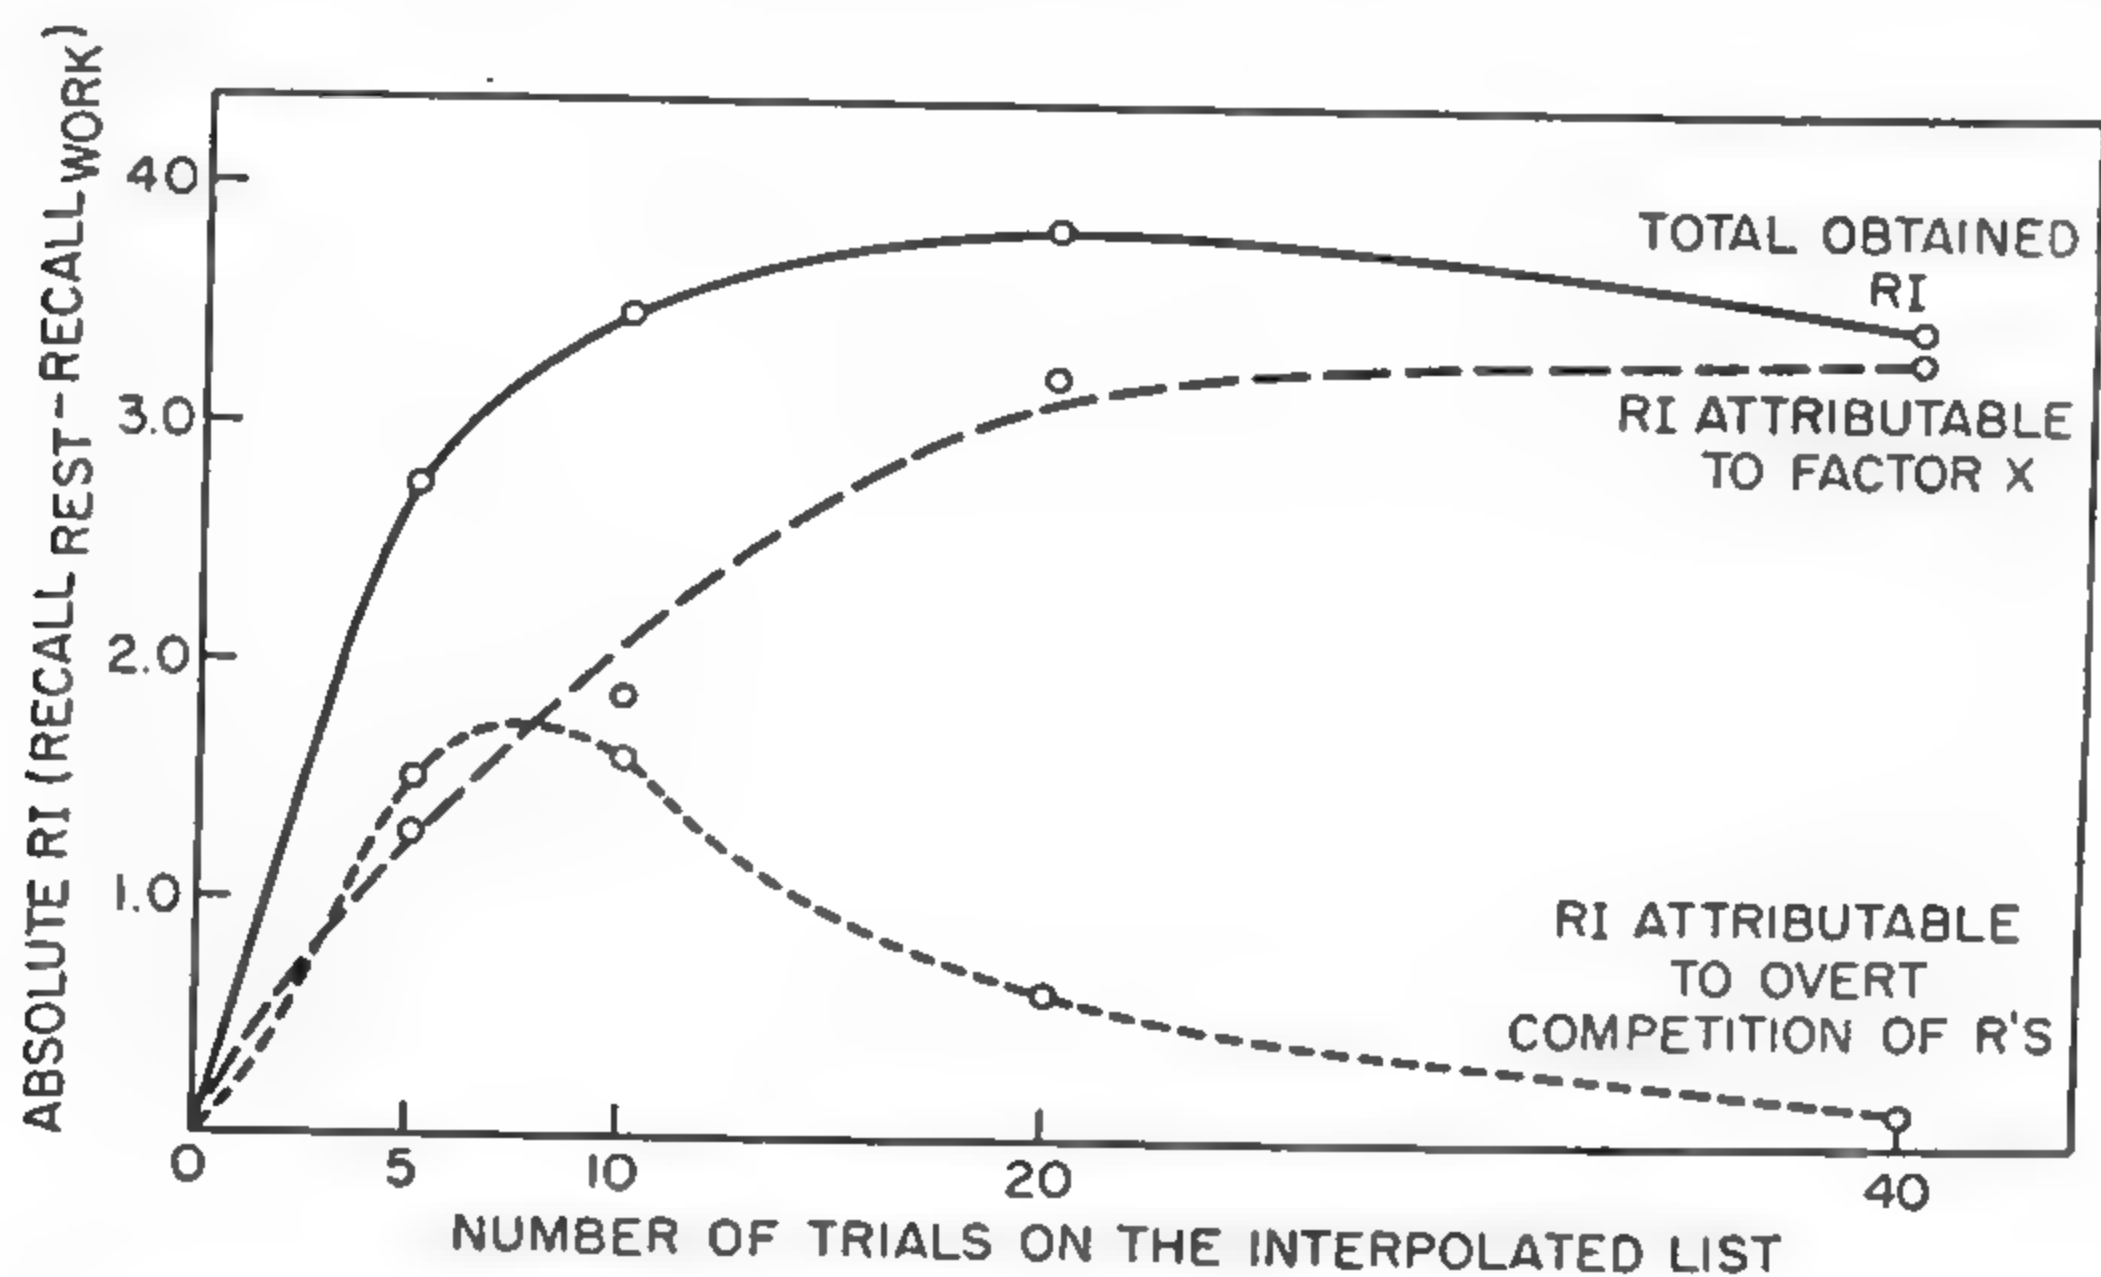


Figure 51. Retroactive inhibition as a function of the number of trials of interpolated learning. The dashed lower curve shows the inhibition attributable to overt intrusion at recall. The upper dashed curve shows inhibition attributable to factor X. (Melton and Irwin, 1940.)

original material. This again is very likely the result of efficient recoding which leads to clear differentiation of the two sets of responses. Unfortunately, total retroactive inhibition continued to increase beyond the point at which intrusions of interpolated responses began to decrease. This situation is illustrated in Figure 51; the curve labeled "RI attributable to factor X" is that portion of the total retroactive inhibition that could not reasonably be accounted for by intrusions from interpolated learning interfering with the recall of original responses. The existence of this factor led Melton to a sweeping revision of the competition theory, and the resulting revision is usually called the "two-factor theory" (note: this two-factor theory is not the same as the two-factor theory of conditioning we considered in early chapters).

The two-factor theory. The competition of responses is clearly one factor in retroactive inhibition. Perhaps, however, other factors are at work in retroactive inhibition, and possibly also in proactive inhibition. One of these seems to be specific to retroactive inhibition. As a matter of fact, a difference between retroactive and proactive inhibition of retention in immediate recall has been advanced as one of the reasons why a second factor in retroactive inhibition is needed. Melton (Melton and Irwin, 1940; Melton and von Lackum, 1941) has suggested that this second factor is the "unlearning" of original responses during interpolated learning.

If interpolated responses intrude at recall of original learning in retroactive inhibition, it seems reasonable to suppose that there would be also intrusions of *original* responses during interpolated learning. And, indeed, there are. When responses from original learning occur during interpolated learning they are incorrect, and hence, it may be argued, are weakened. Thus interpolated learning has the double function of (1) building up a set of competing associations and (2) weakening the set of original associations. In proactive inhibition of retention only competition can work, because the competing responses are learned first, and thus the correct responses cannot be "unlearned" because of practice on the competing responses.

It is extremely difficult to obtain any direct evidence for unlearning, because we do not fully understand how it works. Many experiments, however, give indirect evidence for the existence of unlearning (Melton, cited above; Underwood, 1945, 1948). More recently, Briggs (1954, 1957) has charted the course of responses appropriate to original learning and those appropriate to interpolated learning by means of specially prepared recall tests inserted at certain points during learning. These tests show that the frequency of original responses declines as interpolated learning is continued. In other words original learning is weakened. Furthermore, these responses "spontaneously" recover if a time interval is allowed between interpolated learning and a test for retention of the original responses.

This last observation, which has been frequently made (Underwood, 1948, etc.), has led to the view that the unlearning of original responses during interpolated learning is much like the process of extinction. This comparison explains nothing, but it reminds us that unlearning, like extinction, is in part only temporary in nature. Furthermore, again like extinction, just a little relearning can cause all traces of unlearning to disappear.

Thus, it seems likely that the retroactive inhibition effect has at least two causes, (1) unlearning of the original responses during the learning of the interpolated material, and (2) competition between responses at recall.

Proactive inhibition involves only the second factor. The relative importance of the two factors can be assessed by the fact that there is not much difference between retroactive and proactive inhibition of retention when conditions are otherwise the same and little or no difference when an interval of time is allowed to elapse. Thus competition at recall would seem to be the overwhelmingly more important of the two factors.

The generality of competition. Looking back at Figure 47 near the beginning of this chapter we see that Ebbinghaus's retention curve for nonsense syllables fell off very rapidly after original learning. Twenty-four hours after learning, Ebbinghaus showed a savings of about 35 per cent in relearning the same list. The general picture we get from textbooks on the psychology of learning is that such a result is fairly typical of the forgetting of nonsense syllables. Yet, more than one modern investigator has been puzzled to discover, by using college students as subjects (who are much less practiced at learning nonsense syllables than Ebbinghaus was), that their retention is much better after a comparable time interval. The students in the author's course in experimental psychology, for example, typically show 75 to 90 per cent savings after 24 hours. What accounts for this difference?

Earlier we pointed out that there were two basic designs for the study of retention. In one, the same subjects learned a different list of nonsense syllables for each retention interval; in the other, different subjects learned the same list of nonsense syllables. This accounts for the difference between investigations like Ebbinghaus's which show low retention after a time interval and other studies which show relatively high retention.

Underwood (1957) has shown how this difference is of fundamental importance for a general interference theory of forgetting. He demonstrates that the amount of forgetting of any one set of material is an increasing function of the amount of similar material the subjects have learned in the past. The curve in Figure 52 shows this relationship. Underwood obtained this figure by plotting the data from a large number of investigators; it shows very clearly that the per cent frequency of recall is inversely related to the number of previous lists the subjects have learned.

The importance of this point is twofold. First of all, it clearly shows that the ordinary retention curves, in which no experimental attempt is made to produce forgetting, are determined by the same variables in the proactive and retroactive inhibition experiments. This fact places the competition, or more generally the interference theory of forgetting, on a much firmer foothold.

The second point is even more important. Underwood reminds us that most of the emphasis in the experimental production of forgetting has been on the retroactive inhibition design. But these data on the influence of *previously* learned material upon the retention of more recently learned tasks suggest that the proactive inhibition of retention is really the more important design in the experimental study of forgetting.

To be sure, if the two-factor theory of retroactive inhibition is correct, retroactive inhibition involves a second factor, unlearning, not present in proactive inhibition. We have seen, however, that the unlearning factor is

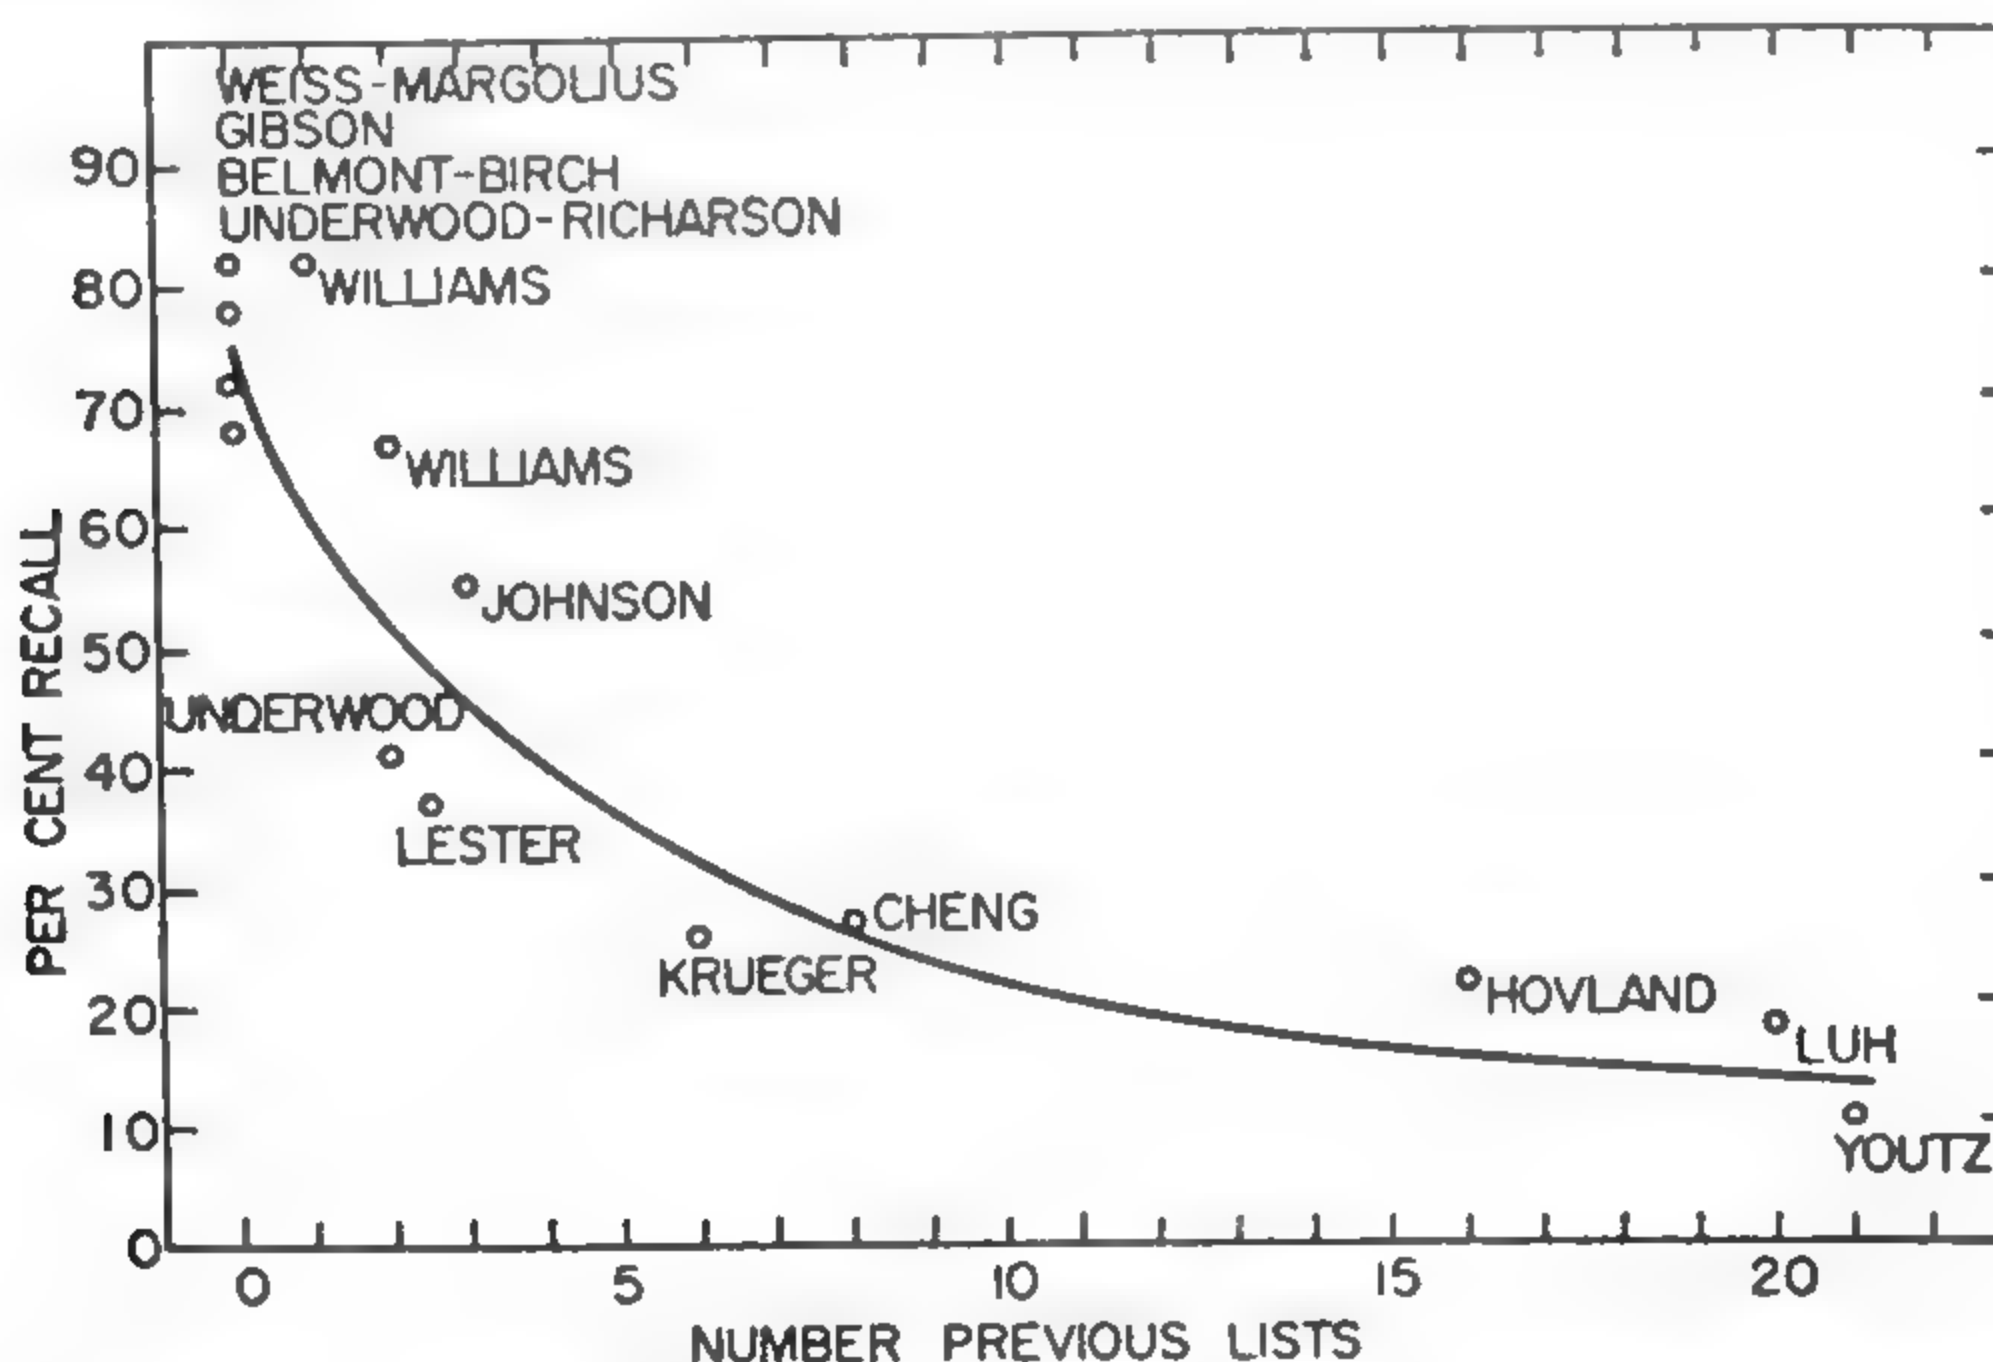


Figure 52. Recall as a function of the number of previous lists of similar material learned by the subjects. This curve was obtained by combining the data of a number of different investigators. The authors of the separate studies are indicated by the names beside the circles. (Underwood, 1957.)

much less important than the competition, or general interference, factor, and furthermore it appears to disappear to a considerable extent with rest.

Therefore, what produces forgetting is, to a considerable extent, simply mutual competition among things we have learned. Fortunately, most of us are safe from the possibility of a complete collapse of our associative processes by the fact that we (1) are continually relearning the things we need in our daily life and (2) we have devised recoding devices for keeping important items free from competition.

Thus I may glance briefly at a set of notes before I go into the classroom to deliver a lecture. This serves to reactivate the coding devices by which I have stored many chunks of information. Just a word or two will enable me to recall a rather complicated theory, and I find as I talk that stories (alas, too familiar) and anecdotes flow automatically.

Direct competition and indirect interference. The method of paired associates provides the ideal model for the competition theory of forgetting. In the classical paired-associate experiment the subject learns two sets of responses to identical stimuli. When the stimuli are presented at recall there is competition between the two sets of responses. In actual practice, in an experiment on retroactive inhibition, competition produced by the alternate set of responses can account for nearly 70 per cent of the errors made by subjects at recall (Deese and Hardman, 1954).

Unfortunately such direct evidence for competition does not always occur in experiments on serial learning by the method of anticipation. If subjects are asked to learn two lists of nonsense syllables by the method of anticipation, retroactive inhibition occurs when recall of the first list is asked for; the second list reduces the tendency to recall the first. Direct competition from the first list, however, accounts for a negligible number of overt errors at recall (about 5 per cent) when recall for the first list comes right after the learning of the second one. What happens in this case is that the learning of the second list causes the subjects to confuse the proper location of items *within* the first (Deese and Hardman, 1954). Thus, while the second list does produce retroactive inhibition, it does so by producing confusion within the original list rather than by directly competing with it.

Thus direct competition after interpolated learning is disappointingly small in serial learning. Apparently, what happens in this case is that the rather weak *organization within* the original list is interfered with. The contextual boundaries between lists, probably established by some primitive coding, are strong enough so that the two lists do not directly interfere with one another. These contextual boundaries are not strong enough, however, to remain intact with the passage of time. If time intervals of a few hours or a few days are interposed between interpolated learning and recall of original learning with serial lists, there is competition at recall from the interpolated list (Deese and Marder, 1957). Thus direct competition occurs if a time interval elapses. The time interval probably weakens the contexts which held the lists apart; we should like to think that this weakening of context is also by interference, but we have no direct evidence for this.

The point to be made is that interference may not always be exhibited in the form of direct competition at the time of recall. These are not grounds, however, for discarding the notion of competition in such cases. We pointed out earlier that competition may depress all "relevant" responses so that irrelevant responses can take over. Thus, the addition of a

second list to the first does not have the effect of producing competition directly, but this addition at fairly low strength (because of the contextual boundaries between the lists) does depress the associations from the first list sufficiently so that their correct order is not maintained.

In conclusion we can say that direct competition is an important determiner of the efficiency of recall. In addition, competition probably affects recall indirectly by competing with the correct responses so that other, irrelevant responses can displace them. Probably the conditions most directly affecting the extent to which competition takes place are similarity among the stimulus conditions associated with particular responses and the degree to which these responses are assimilated (recoded) into the individual's past history of learning.

Repression

The work of Sigmund Freud has been enormously influential in nearly all of the social sciences. The psychology of learning is no exception. The concept of anxiety as a drive, the use of the technique of free association for probing emotional processes, both have their origin directly or indirectly in the theories of Freud. Since Freud touched upon nearly all dynamic processes in human behavior, he had something to say about forgetting. In this section we shall consider one of the most important of the possible forgetting mechanisms suggested by the work of Freud and his followers.

The Concept of Repression

Freud (1925) pointed out that one of the most effective methods an individual can use to protect himself from the conflict between powerful instinctive urges and the restrictive demands of the ego is through the repression of these urges and everything connected with them. This means, first, that the existence of these urges is not allowed to intrude into the individual's conscious life, and secondly, that all mental events associated with these urges are pushed out of consciousness. Thus repression has two forms. In the most fundamental sense repression refers to the blocking of motivation. In a secondary way, however, it refers to an expulsion of certain associations from the conscious life of the individual. Repression requires a constant expenditure of psychic energy, since it must hold down a drive, which itself is expending psychic energy. Thus repression is an active process.

The fact that repression is active does not mean, however, that the individual consciously tries to keep threatening ideas or wishes out of his mental

life. On the contrary, such ideas would have to be admitted before they are suppressed. Freud is arguing that they are not allowed to come into consciousness in the first place. Thus, repression is a dynamic conflict at an unconscious level.

For our purposes, one of the really fundamental aspects of the concept is that anything that is repressed is not lost. It is only submerged below the level at which, under ordinary circumstances, the individual can voluntarily recall it. Furthermore, it does not take relearning to reactivate this material. If the source of conflict is removed, the repressed memory will recover without specific practice. It is this feature of potential recovery that is of critical importance to the role of repression in forgetting.

Experimental Evidence for Repression

Repression is a concept designed to explain certain manifestations of personality, certain kinds of amnesias, the course of an individual's behavior during psychoanalysis, and certain effects under hypnosis. These things, however, only suggest what the role of repression is for associative processes. Can ordinary associations be repressed? The only way to answer this question is to attempt to produce a process similar to repression in the laboratory.

Unfortunately, the experimental evidence on repression is weak. This is partly because repression is not a well-defined concept (we may get involved in an endless quarrel over the meaning of the "ego," for example) and partly because experiments on the production of emotional reactivity, which is necessary to repression, are difficult to perform and interpret. The essential requirements for an experimental demonstration of repression are (1) the establishment of equal degrees of learning of some material in a control and an experimental group of subjects, (2) the imposition of some experience likely to produce repression in the experimental group, (3) indication of repression through forgetting, and (4) recovery from repression after the threat imposed in (2) is removed.

Several experiments meet these requirements. For example, Zeller (1950, 1951) managed to associate the retention of nonsense syllables with the threat of personal failure. Under these conditions subjects failed either to recall or relearn the syllables as well as a control group of subjects. When the threat was removed, recovery of memory occurred. A similar experiment by Aborn (1953), which used somewhat different learning materials and the measurement of "incidental" memory, showed essentially the same results.

These experiments leave many questions unanswered. From the standpoint of the classical theory of repression we may wonder about the role of consciousness. Presumably the subjects in these experiments were motivated to recall the material when under the repression conditions but were simply unable to do so. This is probably the only thing we can say, and in this sense these experiments fall short of psychoanalytic theory (though perhaps the role of consciousness in the theory needs examination). More important questions for our purposes concern the nature of the associations themselves under repression. What happens to the organization of the material? What is the role of associative strength? Can very strong responses or highly coded systems of responses be repressed? Is it possible for us to repress, say, the names of our parents or grandparents? These are all questions which cannot be answered by the experimental data we have, and perhaps they are beyond any of the experimental techniques available to us. Nevertheless, such matters need to be settled before the concept of repression can take the important role it deserves in the psychology of learning.

In Conclusion

We have seen that interaction between responses learned at different times is an important determiner of forgetting. We may well ask, however, whether or not forgetting ever really occurs. If two responses compete at one particular time, the stronger one will win out, but in order for competition to occur, both must be there. Are responses then ever really eliminated? There is, of course, no way to answer so general a question. Nevertheless we can say that the residual effects of something presented in the past can be detected after long intervals of time and under very improbable circumstances. Burtt (1941) describes an experiment with his son, in which he read the boy passages, in Greek, from Sophocles' *Oedipus Tyrannus*. There were three selections of 20 lines each, and for a period of time Burtt read these selections to the boy every day. The child was less than two years old. Six years later, he was required to learn by rote these selections plus some new ones from the same source. It took the boy an average of 435 repetitions to learn the new selections and only 317 repetitions to learn the old ones. Thus, there was savings of better than 25 per cent.

Furthermore, we need to point out that the mechanisms of forgetting discussed in this chapter are not necessarily the only ones. Forgetting is a complex phenomenon, and there are undoubtedly many subtle features of

it that we have not yet begun to explore. We are just beginning to understand the important role of recoding in memory, both in retaining and altering things as they are retained. We need to know much more about how competition and interference interact with recoding. It is probable that once the material is in recoded form, it is relatively resistant to competition, but it may suffer considerably before it is thoroughly recoded. In addition, of course, in its recoded form it may be at the mercy of changes in motivation and attitude.

In short, our past experience is being continuously altered by new experiences and changes in our interests and attitudes. Thus memory is neither a static thing nor the gradual decay of what was once learned. Our past is reshaped for us by what we do now as much as what we do now is directed by our past.

CHAPTER 11

PROBLEM SOLVING AND THINKING

We regard the ability to think as the supreme accomplishment of the evolution of mind. That we should feel this way is understandable, for thinking is the thing in which we human beings undeniably excel. The very fact that we characterize human thinking as a product of the evolution of mind, however, suggests that it has its antecedents in animal mentality. Furthermore, the actual process of thinking and solving problems is learned, and we have already seen that many basic principles of learning are common to man and to higher animals. Consequently, thinking and its analogue in animal behavior must have an intimate connection with the basic principles of learning. In this chapter we shall attempt to outline some of these connections.

PROBLEM SOLVING IN ANIMALS

In the 75 years or so that psychologists have been examining animal behavior in the laboratory, an enormous number of examples of the ability of animals to solve problems has accumulated. Not all of these will help us to understand the nature of problem solving, but a few important ones will show how theories of problem solving have arisen.

A problem exists when an organism is motivated to achieve some goal and there exists no ready means to reach it. Problem solving is the process by which the animal achieves the goal. Since some examples of problem solving merely illustrate the basic principles of instrumental conditioning, problem solving extends down to the most elementary examples of learned instrumental behavior. Other examples of problem solving are quite complicated,

however, and the processes involved are evidently more elaborate than those in simple instrumental conditioning.

Two Classical Viewpoints on Problem Solving

Thorndike's Trial-and-error Problem Solving

The term trial and error applied to the process of problem solving was introduced by C. L. Morgan (1894), but its importance comes from some well-known experiments by Thorndike (1898). Thorndike reported extensive observations on the way in which animals, mostly cats, learned to escape from puzzle boxes.

These puzzle boxes were simply enclosures from which the animal could learn to escape by turning a door latch, pulling a string, or pressing on a treadle attached to the door. Almost every student of psychology knows that Thorndike's cats were very slow and ineffective in solving this problem.

Like most cats, these evidently did not relish being penned in, and they would attempt to escape almost as soon as they were placed in the box. These attempts, however, were neither systematic nor fruitful. After variable amounts of time spent thrashing around, the cats would, quite accidentally as a rule, happen upon the device for opening the door. Usually, as soon as the door was open, the cats could leave. When they were returned to the box, however, they usually began their futile direct escape attempts all over again. It generally took quite a few trials for them to learn to turn directly to the device that opened the door, and the learning was usually an irregular, hit-or-miss affair.

These cats did not appear to be particularly intelligent. They showed little ability to grasp and retain the principle of the puzzle box after having once happened upon it. Thorndike tested a number of other kinds of animals, and while, in general, monkeys and dogs did somewhat better than the cats, he was not greatly impressed with the need for supposing that any very high level of mental activity took place in his animal subjects.

These experiments on the puzzle box do not differ in principle or much in detail from the experiments on the way rats learn to press levers in Skinner boxes. As a matter of fact, Thorndike's experiments led him to state the principle of effect as the basic mechanism in learning. The principle of effect refers to what we described in earlier chapters as the principle of reinforcement in instrumental learning. He pointed out that the essential element in the learning of instrumental behavior is the association of some response with a reinforcement. Gradually, a response so reinforced gains

prepotency over other responses and comes to dominate the behavior of the animal in the testing situation. The behavior of animals is blindly determined by the circumstances of reinforcement. There is little in animal behavior to suggest the need for more complicated processes, or at least these seemed to be the implications of Thorndike's work.

Of course many investigators of animal behavior resented the implications of the pure trial-and-error description of animal problem solving. Both before and after Thorndike's experiments, numerous investigators presented evidence for what they believed to be far more complicated processes in animal learning and problem solving.

The notion that Thorndike was grossly in error and that problem solving in animals is a much more complicated and intelligent process than he supposed finally came to center around the work of W. Köhler on the intelligence of anthropoid apes. We shall briefly review Köhler's work and the implications that have been drawn from it.

Köhler's Insightful Problem Solving

Köhler (1925) studied in great detail the problem-solving behavior of anthropoid apes, particularly chimpanzees. He was critical both of Thorndike's results and methods, and he designed problems for animals which he felt would more adequately probe their mental processes.

Köhler studied the ability of his animals to solve a wide variety of instrumental problems. These were always designed, in contrast to Thorndike's, so that all the elements necessary for solution could be perceived by the animals. There were no trick doors or hidden mechanisms that delivered food. For example, in some of Köhler's best-known observations, he simply suspended the animal's daily ration of food from the roof of the animal's cage. A box was placed some distance from where the food was suspended. The problem for the animal was to perceive the relation between the box and the location of the food and to act accordingly.

In one case, Köhler reports, an animal spent a fruitless day trying to scramble up to the food by various means. When Köhler demonstrated the solution to the animal, however, by placing the box under the food and touching the food while standing on the box, the animal immediately got the solution. There was no trial and error, no contingency between responses and reinforcement. In other examples of the box problem in which the animals more or less arrived at their own solutions, Köhler thought that the animals showed clear evidence of *foresightfully planning behavior in advance*.

Such planning is quite a different matter from trial and error. Since Köhler believed that his animals did solve the problems by foresightful planning, he characterized their solutions as *insightful*. By this he meant that animals did not accidentally happen upon a solution; instead, they saw the relationship between the elements of a problem in such a way that they arrived at an effective solution before they actually engaged in behavior which demonstrated solution.

This means in effect that the animals tried out hypotheses based upon the perceived relationships in the environment. Sometimes, of course, the hypotheses were incorrect, but even an incorrect hypothesis, Köhler argued, is not the same thing as blind, unplanned trial-and-error behavior.

How do we know when insight occurs? It is necessarily an inference from observed behavior. The characteristics of problem solution that led Köhler and other psychologists to identify solutions as insightful are many. Perhaps the most basic is the suddenness of the solution. This is well illustrated in the example given below. In addition, insight is inferred when the solution (if it is correct) persists. Remember that trial-and-error solutions are characterized by irregularity. The animal may take longer to arrive at the solution on the third trial than it did on the second. This is generally not true for insightful solutions. Once the animal gets the solution, it can demonstrate it immediately upon being exposed to the problem. Finally, Köhler claimed, it is possible to infer from the behavior of the animals that they were able to put together objects in the environment that were separate and unrelated. These characteristics of insightful solution are illustrated in an example which is taken from Yerkes (1943).

An example of insightful solution. Yerkes placed an ape in a room with an oblong box that was open at both ends. The box was about 170 centimeters long and very narrow. There was a small door in the top of the box at the middle. This door could be opened by the experimenter, and a banana could be locked inside while the ape watched. A pole was also placed in the room; the pole was just the length of the box. Thus, the ape could get the banana by pushing it out one end of the box with the pole.

Initially, the ape exerted much effort trying to pry open the latch on the door in the middle of the box. Gradually it lost interest in this endeavor and began to show signs of frustration. It even attempted to induce the experimenter to obtain the banana. At one point, however, the animal solved the problem. Its solution came when it began playing with the pole and rolling it around the room. The pole rolled into position alongside

the box. The ape reached into the box with its arm in a vain attempt to reach the banana. It then ran to the opposite end of the box, where it again looked in. Then, suddenly, the ape turned to the pole and instantly pushed it through the box and then ran around to the other side to obtain the banana.

How can this behavior be explained away as an example of instrumental conditioning? Indeed, it is difficult to do so. The ape appeared to see very readily the relationship between the box and the stick. This, by no stretch of the imagination, is similar to the rat's perceiving the relation between pressing the lever and obtaining food in the Skinner box.

It is another matter, however, to say that one therefore needs an entirely new principle of behavior to account for the observations on insightful problem solving. Such a need diminishes further when it becomes apparent that there is a more or less continuous gradation between trial-and-error solution and solution by insight. Furthermore, some observations indicate that a background of simple instrumental learning is necessary for the occurrence of insight. In other words, insight depends upon and grows out of instrumental learning and trial-and-error behavior.

Bridging the Gap between Insight and Trial and Error

The background of insight. Some investigators, not merely content with characterizing problem solving in the higher animals as sometimes insightful, have tried to find the elements in the background of experience necessary for the occurrence of insight.

For example, Birch (1945) presented some chimpanzees with the problem illustrated in Figure 53. As we can see from this drawing, the food was placed within the blade of a hoe, and all that was required of the animal was that it rake the food in. Four out of six animals failed to solve the problem within the 30 minutes allowed. One animal succeeded after it accidentally moved the food into the cage. This, of course, is pure trial and error. Another animal succeeded in what appeared to the investigator as an insightful solution. Thus, despite the simplicity of the problem and the general intellectual superiority of chimpanzees in the animal world, insightful solutions were not plentiful in this example.

Birch went one step further, however. He provided his subjects with sticks to play with. In the course of the stick play, the animals behaved in such a way that Birch inferred that they were using the sticks as functional extensions of their arms. Such learning is probably well described by the

model of instrumental conditioning. The important point, however, is that after a certain amount of play with the sticks, these chimpanzees could solve quite complicated problems, including the hoe problem, which demanded the use of sticks of various sorts.

The implications of this study are clear. Insight depends upon a background of simple instrumental habits. If the animal has not learned these, the probability of insight is extremely low; indeed, solutions to problems are by trial and error. In many instances the occurrence of insight is simply by generalization of lower-order habits. Thus, in some cases at least, insightful problem solving is simply the result of positive transfer.

The extension of trial and error. What about Thorndike's cats? Are

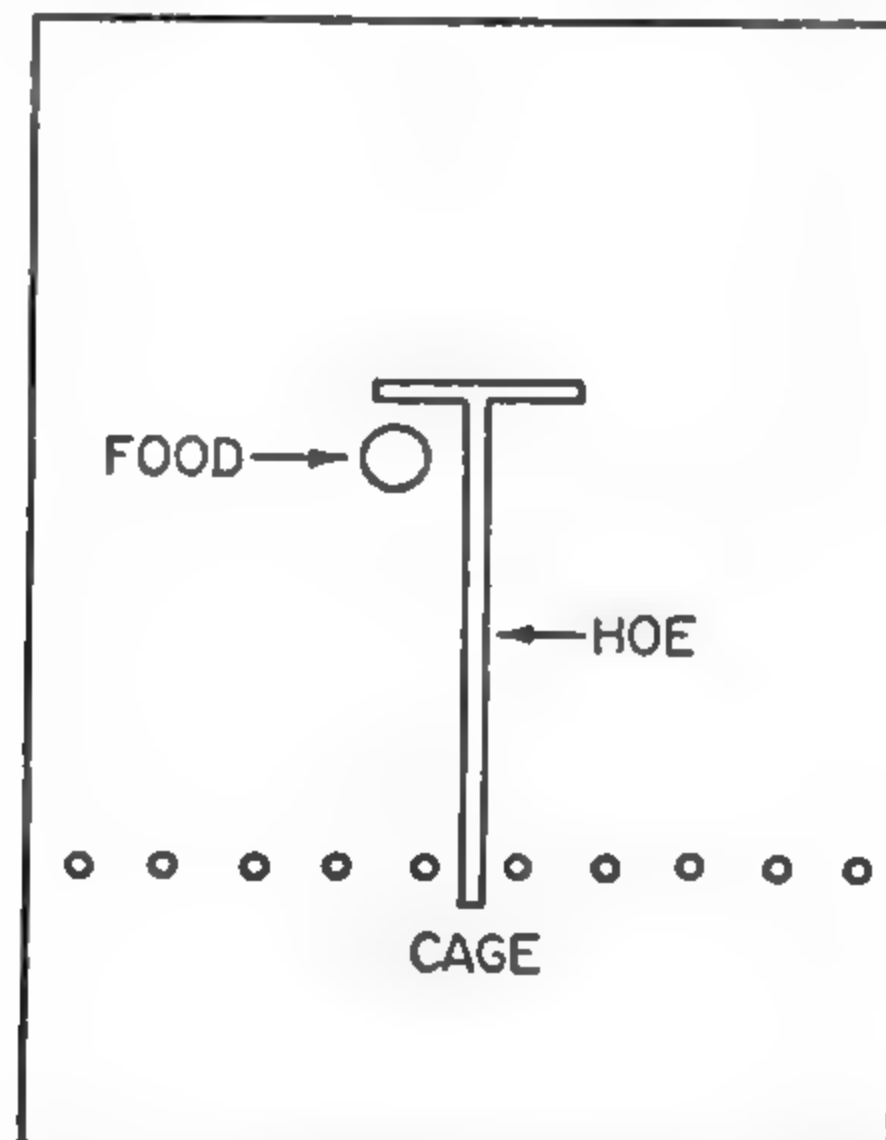


Figure 53. The hoe problem. The chimpanzee can rake in the food, which is out of arm's reach, simply by pulling in the hoe. (After Birch, 1945.)

cats exceptionally stupid animals, or can they too solve problems by insight? This is not a very elegant way of putting the scientific problem, but the fact is that other experimenters have noted considerably more intelligent behavior in cats than did Thorndike. Even in the same situation, the trial-and-error behavior noticed by Thorndike was probably due to inexperience. Adams (1929) repeated Thorndike's experiments and came to the conclusion that Thorndike's cats, in many instances, were simply too excited and untamed to demonstrate anything more than random thrashing about.

Insightful and purposeful solution of problems is not the exclusive property of primates. Indeed, the lowly laboratory rat can solve problems insightfully if conditions are right (Maier, 1929). Furthermore, pure trial-and-error behavior is not unknown among adult human beings. Despite the fact that people use the manipulatory symbolic function of language,

they still indulge in blind trial-and-error behavior when faced with an unusual and difficult problem (Ruger, 1910).

Bridging the gap. Harlow (1951) presents cogent arguments, some of which are like those presented above, to explain why it is not necessary to postulate two basic and distinct processes in problem solving. Trial-and-error behavior is simply the establishment of elementary habits in the absence of stimulus generalization, secondary reinforcement, etc., from previous learning. Such is the case in most experimental studies of animal behavior. The laboratory rat, for example, never in its life encounters anything like the lever in a Skinner box until it is the subject of an experiment on learning. Thorndike's cats had never before been confined in boxes from which it was possible to escape by pulling on strings. Therefore, the correct responses had to be built up from scratch.

In other situations unfortunately all too rare in the psychological laboratory, an organized sequence of training converts a naïve animal into a rather sophisticated one. In many different ways we have had occasion to refer to Harlow's experiments in which he changed monkeys who made choices on successive discrimination problems in a blind and mechanical fashion to monkeys who carefully formed and acted upon hypotheses. These monkeys behaved insightfully, and their behavior came entirely from a rigorous program of training. Their hypotheses were the result of "second-order" habits of observing and testing consequences, and they arose because these are the common features of behavior reinforced in a number of lower-order habits.

Does this mean that problem solving is simply transfer from more elementary habits? It is unlikely, since if this were the case, the only thing to prevent a rat from behaving as intelligently as one of the great apes would be the number and kind of elementary habits it could bring to bear on a given situation. There are genuine species differences in a number of important psychological dimensions, and while we do not know all of them, some investigators have isolated a few that must be very important. Among these are the differences between animals in the capacity to react to stimuli in their absence and to react differently to the second occurrence of a stimulus than to the first. These indicate the ability of animals to react symbolically, and we shall examine next some problems in investigating symbolic behavior in animals.

Symbolic Behavior

It is convenient to divide symbolic behavior into three classes—(1) simple discriminations or classical conditioned responses, (2) reactions to

stimuli which are not in the environment but exist within the organism or are mediated by learned cues, and (3) linguistic reactions in which the organisms recode portions of the environment for action in the absence of the particular cues. It hardly need be said that, so far as we know, this third variety of symbolic behavior is limited to human beings.

Simple discriminations serve to lead the organism from place to place. The formation of symbolic representations in this sense is nothing more than learning in its most basic aspect. The second function of symbolic behavior is to bridge a gap in the absence of external stimulation. This we can call representational symbolic behavior. The higher animals evidently can produce their own symbols, so that they can respond appropriately even after the external stimulus has been removed. The third function, the linguistic function, enables organisms to manipulate the environment symbolically.

It is the second function which is of basic importance to the comparative study of problem solving. If animals solve problems by testing hypotheses, it is essential that they be able to represent to themselves more than the environment present at the moment. To the extent that there are differences among species in this ability, there should also be differences in ability to solve problems insightfully. Consequently, many investigators have examined animals at different positions in the phylogenetic scale for their ability to react to stimuli in their absence. Some evidence on this question is presented below.

The Delayed-reaction Experiments

W. S. Hunter was one of those who realized the importance of examining symbolic action in the problem-solving behavior of animals. Hunter (1924) devised and examined several tests which he thought would force animals to react to stimuli in such a way that the only inference one could reasonably make was that symbolic representation had occurred. Some investigators have disagreed with Hunter in this matter, so it will be necessary for us to examine these tests in some detail. The first is the delayed-reaction test.

There are two basic methods in this test, direct and indirect. In the indirect method the investigator trains the animal to associate a certain stimulus with a reinforcement. For example, Hunter presented animals with three doorways, one of which was lighted. If the animal chose this one, it was reinforced. When Hunter was certain that this association had been established, he would turn the light on, let the animal observe the correct door, and then turn the light off while restraining the animal by means of

a glass barrier. The animal would be restrained for a period of time and then released. If the animal chose the correct door more frequently than chance, Hunter thought there was evidence that it could keep a symbolic representation of the light.

In the direct method, pretraining is not necessary. Generally the direct method works with primates and higher animals. For example, a monkey may be presented with two cups. It watches the experimenter place a raisin under one of the cups. The monkey is restrained for a period of time and then allowed to choose one of the two cups.

Hunter tested a variety of animals by the indirect method and found that there were marked differences from species to species of mammals. For example, in his original experiments (1912), he found that rats could delay a maximum of about 10 seconds, while dogs could delay up to 5 minutes. A systematic ordering of the animals Hunter tested indicated that he had a powerful technique for differentiating an important mental process among the higher animals.

Not everyone was convinced that the delayed-reaction test permitted an unequivocal inference about symbolic representation. Even in Hunter's observations, there were some indications that animals were using continuing external stimulation to bridge the gap. In other words, they were using substitutes for direct symbolic representation. For example, Hunter observed that his rats, when successful at delay, maintained the bodily orientation that they had at the time the light went off for the entire period of delay. Other animals, such as raccoons, did not need to maintain the correct orientation in order to respond. Thus, it seemed likely that the raccoons were more capable of a true symbolic representation than were the rats.

Under other conditions, however, rats can delay up to four minutes without a maintained postural orientation (McCord, 1939; MacCorquodale, 1947). Here, animals were prevented from responding to spatial cues when the correct stimulus was moved to another door. This procedure for preventing response to spatial cues removes certain objections that some psychologists have had to using the delayed response experiment as an indicator of symbolic processes (Seward, 1948). Nevertheless, it is possible to view delayed-reaction experiments of the sort we have been discussing as examples of complex discrimination learning (Seward, 1948). In either case, the solution of the problem represents a high level process in the animals and one that is probably closely linked with the differential ability to solve problems.

The Double-alternation Problem

Another test that has been generally considered to be something that could be solved only by symbolic representation is the double-alternation problem. In the double-alternation maze problem (Hunter, 1928) an animal is required to learn a maze consisting of two alternative pathways. Each pathway describes a square and brings the animal back to the starting point. Figure 54 shows an example of such a temporal maze. The animal must learn to alternate the pathways in a sequence. In the double-alterna-

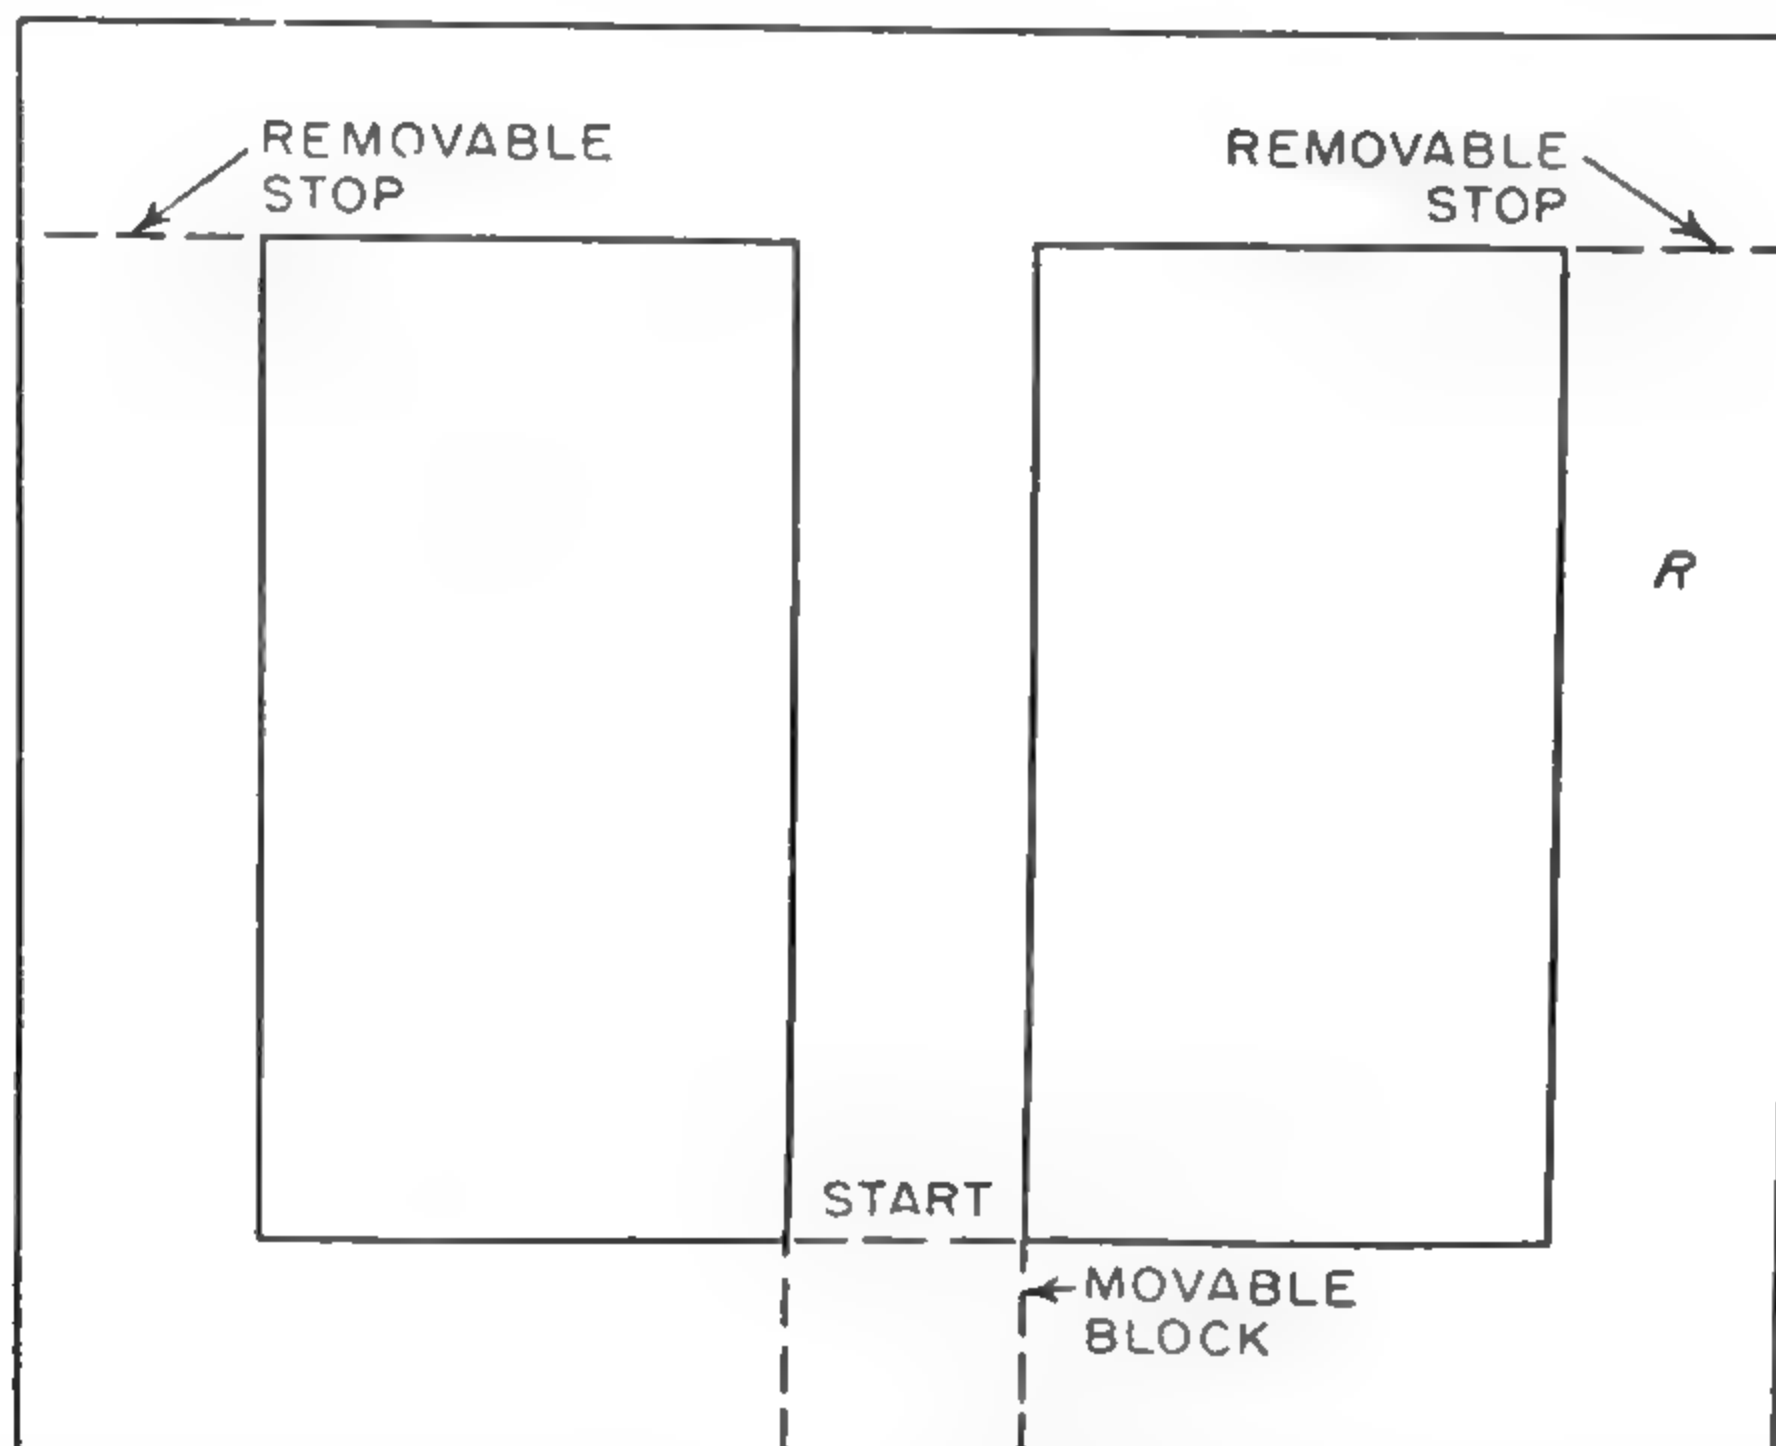


Figure 54. The floor plan of a temporal maze. In the double-alternation problem the animal is required to run around twice in one direction and then twice in the other direction before being rewarded.

tion problem the animal is required to go twice to one side and then twice to the other side in order to receive the reward. The sequence can be extended, so that instead of a series such as RRLL before reward, the animal must learn a sequence such as RLLRRLL.

The rat finds it all but impossible to solve even the simplest sequence (RRLL). Raccoons (Hunter, 1928), cats (Karn, 1938), and other animals can solve the double-alternation problem. The errors, incidentally, in solving such a problem are distributed among the alternatives in much the same way they are in the linear-maze problem (Stewart and Warren, 1957). Preverbal children find the double-alternation problem difficult, but children who can verbalize solve it easily (Hunter and Bartlett, 1948).

It is difficult to see how any animal could solve the double-alternation problem as it is seen in the temporal maze, without some central mechanism that can tell the animal where it is in the sequence. The central mechanism needed is of a fairly low order, considered from the human point of view, since it involves, at the most, the ability to count to two. The double-alternation problem, however, seems to be a more clear-cut case of symbolic behavior in the absence of secondary supporting cues than does the delayed-reaction problem, since it is almost impossible to find grounds for external cues that could help an animal solve it. As Hunter points out, even kinesthetic cues must be ruled out in the temporal maze.¹

The Role of Symbolic Representation

There is something more to insightful problem solving than simply the cumulation of transfer. Some animals, high on the phylogenetic scale, perform in such a way that it is almost impossible to avoid the inference that they are capable of symbolic representation. External stimuli become re-coded in some way so that these animals can react to cues in their absence or react in such a way that they can keep separate the first and second occurrences of a particular external cue.

Symbolic representation can have many different functions in problem solving. For example, it probably plays an important role in multiple-choice problems which require an organism to remember and reject the alternatives already tried and found wanting (see Miller and Frick, 1949). It must be fundamental in helping animals to see relationships among external events. In the example of insightful problem solving described a few pages back, symbolic representation of the pole while the animal was peering into the oblong box would almost seem necessary to an insightful solution to the problem.

Unfortunately, the implications of some of the basic properties of symbolic representation have never been fully explored. Physiological psychologists who are interested in investigating some of the localized functions of the central nervous system have made extensive use of the tests we have mentioned. Few investigators, however, have tried to follow some of the implications of these tests for problem solving generally. All we can do at this stage is to point to the tests themselves and to the differential ability among animals for problem solving.

We know that there are marked species differences in the ability to trans-

¹ Although in the double alternation of lever pressing, which the rat can solve, kinesthetic cues could conceivably work (Schlosberg and Katz, 1943).

fer complex discrimination problems (Harlow, 1949; Cotterman, Meyer, and Wickens, 1956, etc.). These probably also reflect underlying abilities to use symbolic representation. It is intriguing, though perhaps fruitless, to speculate on the nature of the differences in representational factors among different animals. Some investigators (Stewart and Warren, 1957) have suggested that in many tests of symbolic representation, the differences are quantitative rather than qualitative. In the normal adult human being, though, there are clear qualitative differences.

HUMAN PROBLEM SOLVING

We are sure that normal adult human beings can solve problems symbolically, because they possess highly developed and organized linguistic tools for doing so. Therefore, there is no question whatever about the normal adult's ability to solve problems insightfully. He may not always do so although he is capable, and, indeed, a knowledge of this capability is what leads us to suspect an analogous process among animals based upon a much more rudimentary symbolic process.

For many years, Piaget and his associates have been studying the development of intellectual processes in children. These studies fill out in rich detail the growth of the processes of insight and concept formation in developing children. A great deal has been learned about children's concepts of number, space, time, and other fundamental matters from these studies. In addition, these investigators have studied the actual process of problem solving. For example, Inhelder and Piaget (1955) point out that young children solve problems by finding similarities and differences among *things*; their approach to problems is concrete. Adolescents, on the other hand, may develop formal propositions; they try to formulate or find abstract laws for solving whole classes of problems irrespective of their specific content. Although it is on a different level and involves quite different processes, this description of progression from concrete to abstract problem solving is highly reminiscent of the formation of learning sets in Harlow's monkeys.

An account of such factors, however, is more properly the subject matter of developmental psychology. Our task at the moment is to find out what we can about the relationship between associative factors, learning, and the solution of problems. Therefore, we shall turn to the experimental literature aimed at the analysis of problem solving and the relationship between this process and mechanisms of learning.

Experimental Studies of Human Problem Solving

The Formation of Hypotheses

Hypotheses are not random. When a normal adult is faced with a problem to solve, his attempts at solution are not random samples of the possible behavior in which he might engage. For that matter, the cat in the puzzle box does not really behave randomly either. Certain classes of responses have preference over others for both man and cat. The behavior of the cat is more likely to be characterized as trial-and-error behavior simply because, in all likelihood, it will be less systematic and more random than the ideal strategy for solving a particular problem. As we shall see shortly, men do not always or, indeed, frequently employ an ideal strategy, but their behavior is more likely to be systematic than the behavior of the cat or monkey. In this sense, man is less likely to engage in pure trial-and-error behavior than are any animals.

There are many ways in which the behavior of men and that of animals differ qualitatively during problem solving. An animal, for example, is much more likely than the normal, unruffled adult to repeat over and over again a response which has already been tried and found not to lead to solution. The range of possible attempts to solve a particular problem will, in general, be less in animals than in man. Finally, man, by his ability to use language, can much more effectively recode the problem into more familiar terms or he can exhaust systematically the possibilities for finding a solution.

Thus, the hypotheses men use in solving problems are not random selections of the ones they might use. They tend to be determined by rather systematic factors. Cofer (1954) points out that some of the most important determining factors in human problem solving are mediating verbal responses.² These are responses that the individual is not likely to make overtly; they are implicit verbal responses of the sort we usually mean when we talk about thinking. The important verbal mediating responses which lead to attempted solutions of a problem may not be directly elicited by the problem situation but may be only indirectly elicited by virtue of association with responses directly produced by the situation. In the next few pages we shall examine some of the ways in which verbal mediating processes influence problem solving.

² Also see Osgood (1953) for a detailed account of the theoretical role of mediational processes in learning and thought.

Verbal mediating responses and direction. The way in which verbal mediating responses can give direction to hypotheses in problem solving is illustrated in an experiment by Judson and Cofer (1956). In this study, subjects were given groups of four words and told to indicate the one word not related to the other three. Suppose, for example, that ADD, SUBTRACT, MULTIPLY, PER CENT were given. This item is easy. It is obvious that PER CENT is the word that does not belong, since it is the one word that does not indicate an arithmetic operation. Suppose, however, that the words were ADD, SUBTRACT, MULTIPLY, INCREASE. Here the answer is not so easy. To be sure, the same three arithmetic operations are present and INCREASE is not an arithmetic operation. But ADD, MULTIPLY, INCREASE all belong to a class of words indicating increasing magnitude, while SUBTRACT belongs to a class of words indicating decreasing magnitude. Perhaps SUBTRACT is the word that does not belong.

Judson and Cofer administered many such ambiguous items to subjects. The object was to find out which ambiguous word the subjects would accept as belonging with the other words. One of the things they discovered was that the order of the words was important. Thus SUBTRACT, INCREASE, MULTIPLY, ADD would more likely lead to INCREASE being rejected, whereas MULTIPLY, INCREASE, ADD, SUBTRACT would more likely lead to SUBTRACT as the rejected word.

Secondly, subjects were likely to accept ambiguous words that conformed to their previous habits. For example, in problems like PRAYER, TEMPLE, CATHEDRAL, SKYSCRAPER, subjects with strong religious interests were more likely to accept the word PRAYER and reject SKYSCRAPER than subjects with little religious interest. Thus it seemed clear to Judson and Cofer that the direction chosen by the subject was the result of activation of mediating responses which recoded the material in a particular way. The probability of a particular mediation response is determined in turn by its relative associative strength. This point will become quite clear in later examples.

Fixation in Problem Solving

Fixation by mediating responses of function. If a particular problem elicits a strong set of mediating responses which point to a particular solution, and if that solution is incorrect, we have a case of fixation without adequate solution of the problem. The subject becomes fixated on the wrong response and cannot switch to the right one. Frequently problems are pre-

sented to us in such a way that they do elicit the wrong mediating responses, and great difficulty with the problem results.

This point is illustrated by some of Duncker's (1945) well-known observations on problem solving. Duncker gave students a variety of simple problems to solve. One of them was as follows: The subjects were required to mount three small candles on a door (ostensibly for experiments on vision). The materials the subjects needed for solution were scattered in confusion on a table and mixed in with a number of other objects. The crucial items were tacks and some cardboard boxes similar to small match boxes. The solution required the subjects to tack the boxes to the door, and then to melt wax on the boxes in order to hold the candles.

This problem was presented to the subjects in two ways. The difference was small but of critical importance. In one condition the boxes were filled with the experimental material—tacks were in one box, candles in another, and the matches in the third. In the other condition the tacks, candles, and matches were placed loose on the table and the boxes were empty.

What difference did these conditions make? Those subjects who were presented the problem with the boxes empty all solved it, while only about half the subjects were successful when the boxes were filled. What was responsible for this difference? It is very easy to understand if we consider the possible mediating responses that these two versions of the problem elicited. When the boxes were filled, they were more likely to elicit a mediating response such as "container." When the boxes were empty, the mediating response "container" was likely to occur also, but not nearly so strongly. The result was that other mediating responses, among them perhaps "platform," were able to occur. The filled boxes thus produced a fixated function which interfered with the proper solution of the problem.

Duncker studied a number of different problems like the one we have just examined. His experiments were carried out with a good deal of ingenuity, but, unfortunately, they leave much to be desired, since the number of subjects was very small and the experiments were performed under rather informal conditions. Other experimenters who have worked with these or similar problems have achieved much the same results, however, so this functional fixation, as Duncker called it, seems to be a genuine effect (see Adamson, 1952).

One important point comes out of subsequent experimental work. If time is allowed between the exposure which sets off incorrect fixation and the actual work on the problem, the fixation disappears (Adamson and Taylor,

1954). Thus, it is clear that whatever kind of verbal mediating habits are responsible for the fixation, they behave much as other kinds of verbal habits.

Other experimental work has indicated the effect of verbal responses suggested to the subject upon the direction problem solving takes (Maier, 1930, 1931). In one experiment, merely teaching the subjects lists of words which could arouse associations pertinent to solution increased the frequency of relevant solutions (Judson, Cofer, and Gelfand, 1956). Cofer points out that such activation of appropriate solutions probably occurs through mediated generalization (Cofer and Foley, 1942). For example, if the word "rope" is associated with "swing," subjects are more likely to solve a problem which involves getting two strings tied together by swinging one to the other by means of a pendulum bob. If, on the other hand, "rope" is associated with "hemp" and "pendulum" with "clock," the pendulum swinging solution is less likely to occur.

Riddles, puzzles, and similar brain teasers are based on the probability that a particular way of stating a problem will elicit a very strong association that prevents the appearance of a mediating response leading to the correct solution. This same sort of thing often makes students complain about trick questions on examinations.

Fixation by method. In the examples we have discussed thus far, direction of solution has been seen to be determined by verbal mediating processes either set off by the problem situation itself or suggested by verbal associations given to the subject. Fixation, however, can occur by other means as well. Probably one of the most important is the pattern of reinforcement for particular kinds of solutions. This kind of fixation is illustrated by some experiments by Luchins (1942) on a well-known kind of problem.

Luchins asked a group of subjects to solve a number of problems all of the same type. Here is an example: Given a 3-quart jar, a 21-quart jar, and a 127-quart jar, how would one measure out exactly 100 quarts? The solution is to subtract 21 quarts from the 127 quarts, which makes 106 quarts, and then to subtract 3 quarts twice. The first problems given to the subjects always required the use of all three measures for an efficient solution. After the subjects had solved six such problems, others were introduced which could be more efficiently solved with only two of the measures, though all three measures could be used. Not surprisingly, perhaps, the subjects kept on working through these simpler problems by the more complicated method. Thus the habit built up by the earlier problems persisted,

even though it no longer provided the best method. When the subjects were instructed to write "don't be blind" on their papers immediately after the sixth problem, a much smaller number persisted in the more difficult method.

Such persistence of habits no longer most efficient for solving particular problems has been characterized as rigidity, and a number of investigators have experimental evidence to indicate that this kind of rigidity is modified by the scheduling of problems and hence of the habits that are reinforced. For example, on a conceptual card-sorting task, Schroder and Rotter (1952) demonstrated that a training series requiring frequent shifts in set resulted in much greater flexibility than one that required no shifts.

Maltzman and Morrisett (1952; 1953) performed a series of experiments in which they produced various fixations on the solution of anagram (scrambled words) problems. They pointed out that the effects produced by fixation are easily interpreted as the result of simple habits. As a matter of fact, they were able to derive and verify certain predictions about fixation from the assumption that the strength of the habits underlying fixation is a simple negatively accelerated (exponential) function of the number of previous elicitations of that habit. Remember from the discussion in Chapter 2 that such a simple negatively accelerated curve is the most basic and elementary kind of learning curve. Thus, there is every reason to believe that the habits responsible for fixation are determined by all the basic principles of learning.

Hypotheses in human problem solving are determined by the arousal of verbal mediating responses, and these are aroused by and arouse in turn simple associations. These associations may or may not lead to the correct solution of the problem, and to a certain extent the individual will try to direct or restrict associations to areas which have some probability of leading to a solution.

The habits that are aroused by the verbal mediating processes vary in strength. The stronger habits are those that have occurred more frequently in the individual's past history, and these have a greater probability of occurring in any problem situation than weaker habits. Furthermore, they depend upon the stimuli with which they have been associated in the past. For example, a pendulum bob has a very low probability of eliciting responses appropriate to a hammer (using the bob to drive a nail into the wall). Yet, when the appropriate conditions occur, subjects can vary their mediated responses to the bob in such a way that it occurs to them to use it as a hammer (Duncker, 1945).

When a secondary or indirect use of a familiar object occurs, it is probably because of either direct primary-stimulus generalization (which would be greater, say, if the pendulum bob looked like a hammer head or because of mediated generalization. Mediated generalization occurs along some learned, or conditioned, aspect of stimuli and is characteristic of the relationships between words in a language (Cofer and Foley, 1942). Thus, the word "vane" might give rise to its homonym "vain," which in turn might arouse "fruitless" and this could lead to "barren" and then to "dull" or "stupid." By an elaborate intertwining of learned and unlearned generalization between words (considering words as stimuli rather than items of behavior), a great flexibility is added to human behavior. This must be of fundamental importance in human problem solving.

Furthermore, many factors which are not primarily associative determine the likelihood of particular solutions or outcomes in problem solving. These are so-called dynamic factors in behavior, such as motivation, attitude, and personality structure. There is a fairly large literature on the relationship between these factors and problem solving. For example, Thistlethwaite (1950) was able to demonstrate that students with particular biases concerning racial issues make more errors drawing inferences from statements with racial themes than from statements with neutral themes. The greatest number of errors in drawing such inferences does not necessarily occur with emotional items that are ambiguous, but rather with those in which the individual can see a possible answer compatible with his views.

Since such attitudes and, to a certain extent, motives are determined by associative factors, simple associative factors have great importance in determining the nature and direction of the attempts people make to solve problems.

Other Aspects of Problem Solving

We have had nothing to say about two very fundamental aspects of problem solving. What is the nature of understanding? How does an individual recognize and pick out the correct solution from a set of many available possibilities? The discussion of these questions, while basic to an understanding of human thought and problem solving, would carry us away from the central issues in the psychology of learning. We mention them at this point to make clear that we do not solve all the problems of thought by discussing the habit structure of verbal mediating responses. Individual differences in ability to solve problems and to think creatively cannot be adequately described in terms of simple habits, but an understanding of the

role of habit, transfer, and generalization in problem solving gives us a grasp of certain fundamentals of the process.

The Learning of Concepts

The world presented through our senses is a vast jumbled confusion of different sensations. We are able to deal with it only by cutting it down to the size of our own mental processes. The primary way we do this is by setting up equivalences and identities among separate parts of our experience. In short, we categorize and assign names to the categories. Thus, the deep maroon color of the book on my desk and the tomato-colored stripes in my tie both are characterized as "red." This process of categorizing so that all the infinite variety of the external world may be dealt with by our mental processes and language, is one of the most essential elements in human thinking. We call this process *concept attainment*, or *concept formation*. In the next few pages we shall examine some of the experimental literature on concept attainment and see how special problems in the organization of concepts are related to the basic principles of learning.

Forming Concepts

Categorization. Most of us are so familiar with the conceptual categorization of our world that we do not stop to think about the process of arriving at categories. We may even naïvely believe that our conceptual categories have independent existence. The habit of accepting the reality of conceptual categories may be so firmly established that it will interfere with scientific progress and thinking. Only in the past few years have systematic biologists generally come to the view that "species" of animals are not fixed classes into which animals do or do not belong, but rather that "species" are simply convenient classes into which we can group animals with similar characteristics. Nature is often continuous and will not fit the boundaries of our categories. Thus, frequently there is a continuous gradation in essential characteristics from one species of animals to another.

Every college student at one time or another has run into a homely example of the conflict between categories and the continuity of things. Registrars and deans demand that instructors categorize students in classes usually labeled A, B, C, D, and F. There may be only the faintest difference between the student with the highest B and the lowest A, but because human society and thinking cannot function with continuously graded events, and must resort to categorization, we are required to place these students in different classes.

We form classes by isolating attributes of things which are to be classed as identical or equivalent. There are two ways of doing this. We may take a heterogeneous collection of things and examine them for common attributes so that we can sort them into classes, or we may come to examine a number of things with some preconceived notion about the classes into which they ought to be placed. In either case, we assign names to the classes (thus recoding our experience).

Our categorization is usually done in accordance with social reinforcement. Thus, children learn at an early age to discriminate between dogs, cats, and rabbits, though they may make mistakes during the early stage of learning to assign the proper names to the individual animals.

Learning conceptual categories. Osgood (1953) reminds us that the formation of concepts involves verbal mediating responses which permit us to recode the continuity of our experience. The classical experimental literature on the formation of concepts has dealt with the ways in which such verbal mediating responses are established.

For example, Hull (1920) studied the way in which students learned to find identity relationships between stimuli possessing common elements. He presented subjects with Chinese characters paired in a certain way with English words. Chinese characters are compounded of certain elements, called radicals. These may vary in position or size within the character, but whenever they are present it means that the character has something in common with other characters which also contain the same radical. They are roughly analogous to certain syllables (mainly suffixes) in English. Thus, the words "repeat" and "return" have in common a syllable (actually, a bounded morpheme) which carries the connotation of coming back over the same event. In Hull's experiment, whenever a certain radical occurred it was always combined with a certain English word. The problem for the subjects was to recognize and associate the radical with the appropriate English word.

As might be supposed, such learning is much the same as any other kind of paired-associate learning. There were, however, some cases in which subjects obviously had learned to associate the proper English word with the proper radical, but these subjects could not recognize the radical. Thus they were reacting to something they could not verbalize. They had learned the correct pairing of the Chinese characters with English words, but they had not acquired a mediated recognition of this.

Other experimenters (Heidbreder, 1924; Smoke, 1932) have studied the role of hypothesis formation in the discovery of concepts. Smoke (1932)

reports that his subjects systematically tested and rejected a number of hypotheses—thus making use of transfer—until they finally came upon one which gave them enough correct responses to be satisfying. Heidbreder (1924) emphasized the fact that the learner does not always have to be actively engaged in hypothesis formation in order to arrive at a solution of the problem. In especially difficult problems some subjects engaged in what Heidbreder called “spectator” behavior; these subjects would resign themselves to responses which were more or less random and would spend their time in passive observation of the material presented to them. Under these conditions, subjects eventually solved the problems. Thus a period of quiescent observation may be just as valuable as a period of active hypothesis formation.

The attainment of specific concepts. Despite the fact that concept attainment does not differ qualitatively from other kinds of learning there are differences in the ease with which various concepts are learned.

Heidbreder (1946a, 1946b, 1947) had subjects learn to attach labels to the concepts illustrated in Figure 55. From the pictures we can see that some of the concepts were highly concrete (faces, buildings, etc.), and others were unfamiliar and abstract (“nonsense” forms, numbers, etc.). Heidbreder found that there was a regular order in the attainment of concepts. The concrete-object concepts were the easiest to learn, the nonsense form concepts of intermediate difficulty, and the abstract number concepts most difficult. Heidbreder attributed this to a special factor she called the *hierarchy of dominance*, which is due to the objectlike character of perception. Thus, the concept of “face” is easy to attain, since the instances of it are perceptual entities—in a word, faces. Numbers are difficult to attain because the instances of them are embodied in concrete objects rather than being concrete objects themselves. Thus, the concept of “two” might be exemplified by two spoons or two rabbits.

A subsequent series of experiments (Dattman and Israel, 1951) did not confirm Heidbreder because different ways of embodying concepts were used. When appropriate instances of numbers or abstract figures are used, the order of attainment of concepts disappears. This finding makes it quite evident that the ease of attainment of concepts most certainly depends upon the stimuli and the relationships between stimuli used to exemplify the concepts.

It turns out that there is another way to describe the differences in ease with which concepts may be attained, and this may be more fundamental than Heidbreder’s. An experiment by Baum (1954) clearly shows that the

with three, since frequently the stimulus embodying the concept of "sixness" was given in two rows of three objects each. Whereas the subjects found the concept of "faceness" quite easy; they would only confuse it with the concept of hat. Thus stimulus generalization (or, more accurately in some instances, mediated generalization) produces confusion between the examples themselves, which leads to slower learning of some concepts.

Again, as we might expect, since concept attainment seems to obey all the rules for other kinds of learning, the rate of attainment of a concept is determined by discriminability among the instances of different concepts. Therefore, argues Baum, it is not necessary to say that an abstract concept is more difficult simply because it has less perceptual thing-quality about it (as Heidbreder seems to imply), but rather because an abstract concept is generally embodied in specific instances that are less easy to discriminate from those of other concepts. Thus a child has less difficulty learning the concept of "tree" than that of "weed," because trees are more discriminable from other growing things.

It is difficult to say whether or not one could ever find a situation which clearly made a difference between Baum's reasons for the variations in ease of attainment of concepts and Heidbreder's, for nearly always, concepts which are easily seen as intact perceptual objects are easier to discriminate than concepts which are not concrete things. Indeed, it may be that Baum and Heidbreder are putting forward much the same idea. Baum's version seems to be the more general one, however, and it relies on ideas derived from more elementary problems in learning, whereas Heidbreder's notion is a specific hypothesis advanced to account only for the data on concept attainment experiments.

Positive and negative instances. In learning about concepts we may profit from two kinds of examples, positive and negative. Positive examples are instances of the concept and include the essential characteristics. Negative examples are instances of what the concept is not and may lack one or more of the necessary characteristics. Thus a physician may diagnose a particular disease if a high fever, spots, and an erratic pulse are present. If, however, only the high fever and erratic pulse are present and not the spots, it is not an example of this particular disease.

In some early experiments by Smoke (1932) the results suggested that negative examples did not materially contribute to the learning of the concept. In other words, subjects learn by finding out what a concept is, not by what it is not. Hovland (1952), however, pointed out that there is generally less information in negative than in positive examples, and it is not clear

whether the differences obtained by Smoke are because of the smaller amount of information or because people just have difficulty in assimilating negative examples. Therefore, Hovland and Weiss (1953) performed a series of experiments in which the amount of information in positive and negative instances was controlled. It turned out that subjects do better with positive examples even under these conditions, though some subjects could learn entirely on the basis of negative examples.

There are many reasons why positive rather than negative examples lead more readily to concept attainment. For one thing, the required characteristics are directly presented, whereas with negative examples, if the subject is to learn what the concept is, he must put together all of the things a concept is not.

In everyday life, concepts may be almost impossible to establish by negative instances. This is because the ways in which things differ are so enormous that a single negative instance scarcely reduces uncertainty at all. Try to describe a simple object such as a book by mentioning all of the things it is not. It is possible in the experiments by Hovland and Weiss in which negative instances were designed to tell as much as positive instances, that the subjects did not do as well with the negative ones simply because they were not used to dealing with concepts in this manner. A machine could be constructed which would solve the Hovland and Weiss problems just as easily either way. It is clear, however, that human beings have great difficulty with negative examples even when the information presented is considerable. The fact is, though, that when forced to it people can arrive at concepts through negative instances.

Furthermore, as Bruner, Goodnow, and Austin (1956) point out, many concepts must be arrived at largely through negative instances. This is characteristic of the attainment of *disjunctive concepts*. Disjunctive concepts are those in which one *or* another attribute may be used to identify the concept. These authors give as an example the difficulty in determining substances responsible for an allergic reaction in an individual. A food allergist must start by finding out what his patient is *not* allergic to and from there try to arrive at a class of allergy-producing substances by many negative examples and perhaps a few positive ones.

Strategies in Concept Formation

People do not arrive at hypotheses in problem solving by random selection. Rather they employ strategies, though they need not necessarily be aware of committing themselves to particular plans. Strategies are arrived

at by deciding which class of responses one wishes to try in solving a problem or attaining a concept and by deciding in what order these responses are to occur.

For example, I forget my keys one morning, and I must borrow a set from the janitor to get into my office. The keys he gives me are all alike, so I must decide how to find the one that fits. One strategy would be to try the keys in order. Another would be to examine the keys for a clue, such as code numbers. If I am harried and out of sorts, I may simply jam the keys into the lock more or less at random. The last is a strategy too. The problem of strategy is an important one, because we know that there must be a best approach for every problem. Some investigators have designed problems in such a way that they know the best strategy in advance, and these experimenters have tried to find out if subjects use the best strategy, or if not, which one they do use.

Bruner, Goodnow, and Austin (1956) have made use of some special techniques for studying strategies in concept attainment. In a typical experiment, a subject would be presented with a set of cards. These cards varied in four ways: (1) they had different figures on them; (2) they were of different colors; (3) there were different numbers of figures on different cards; (4) they had different numbers of borders around them. Each of these "attributes" had three different values; thus there were three kinds of figures, crosses, circles, and squares. All 81 possible cards in this set are illustrated in Figure 56.

The experimenter would form a concept built on selected "values" from these attributes. For example, the concept might be green circles with two borders. Thus all cards with both green circles and two borders (conjunctive concept) would be correct. The subject's task was to find out what the concept was by choosing a card and asking the experimenter whether this was an example of the concept or not. What strategies do people follow in choosing the cards?

Kinds of strategies. First of all, there is a logically best strategy. In the problem we outlined there are 255 possible concepts. A first *positive* card eliminates 240 of these, and there is a similar value for all the positive and negative instances thereafter. A logical computing machine could keep track of all of the possibilities logically eliminated and solve the problem in the minimal number of trials. A little reflection will make it clear, however, that such a strategy places an almost impossible burden on the human memory and the capacity for assimilating information already obtained. Thus it is very unlikely that anyone would use this completely logical ap-

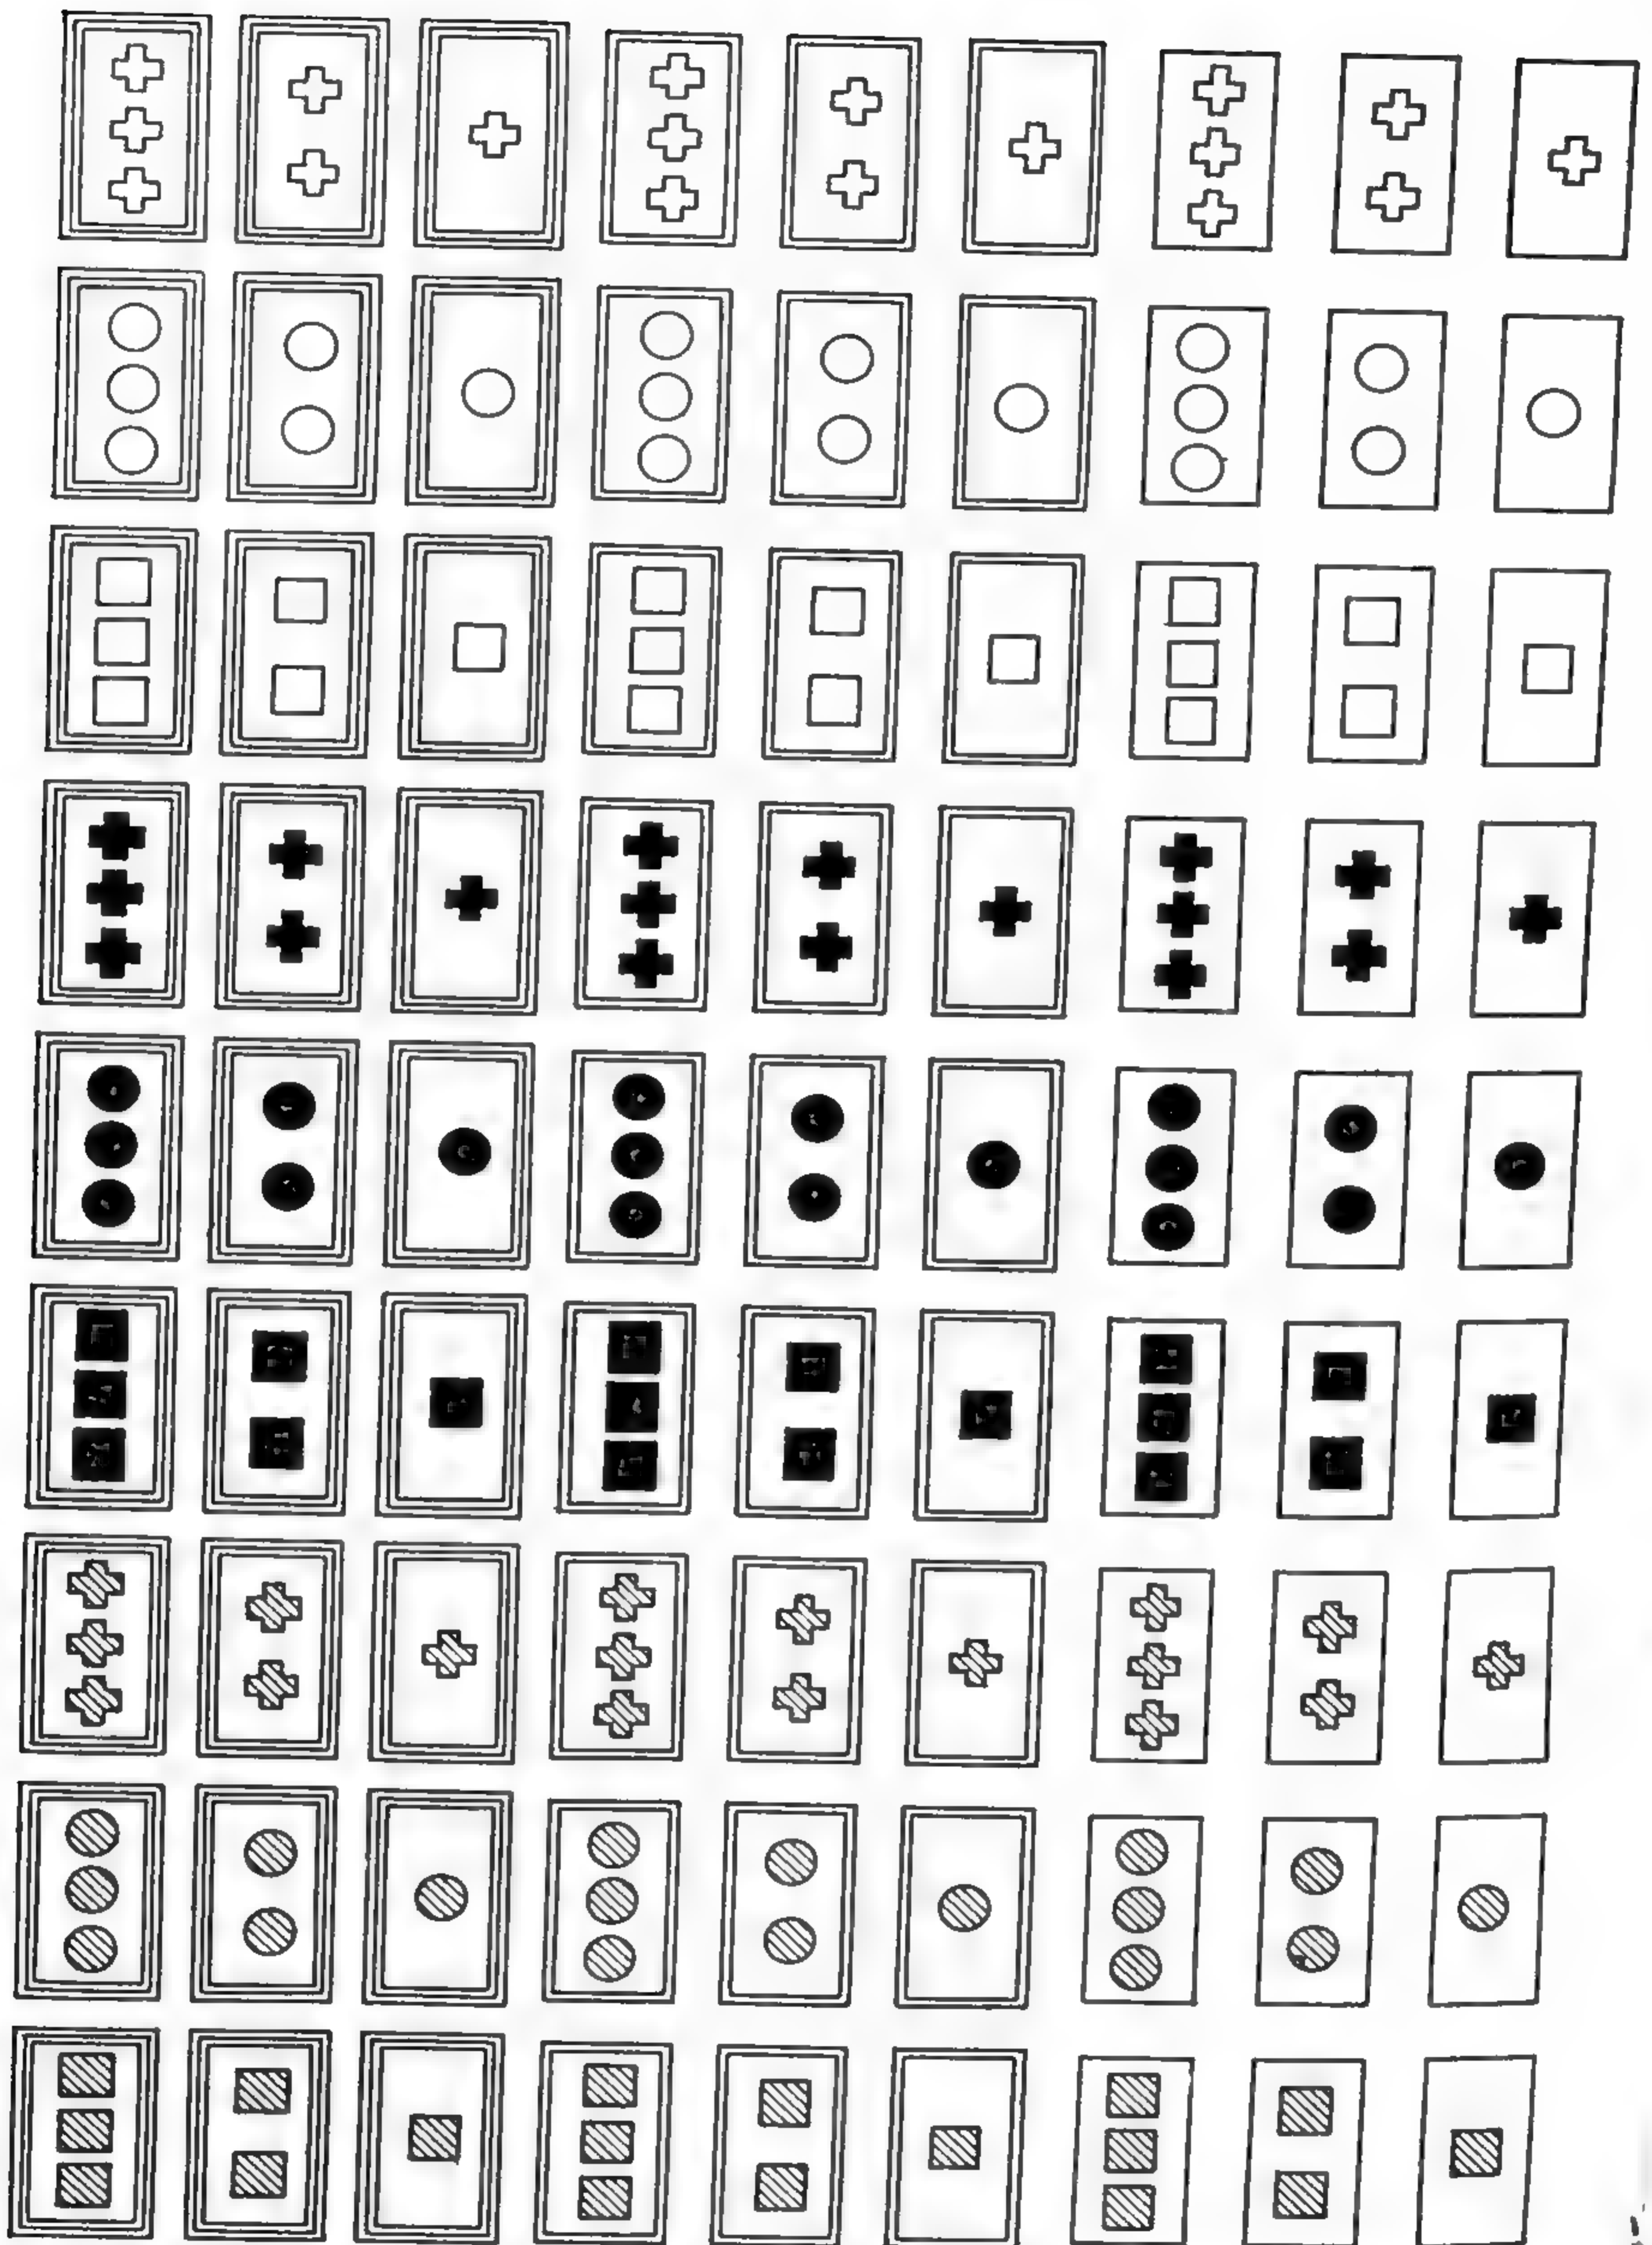


Figure 56. The instances used by Bruner, Goodnow, and Austin to illustrate attributes. The forms varied in (1) shape, (2) color, and (3) number of borders. There were three values for each attribute (e.g., color was either green, red, or black). (Bruner, Goodnow, and Austin, 1956.)

proach. Some subjects, however, do make a reasonable approximation to the logical strategy.³

Another strategy the investigators called *conservative focusing*, and it consists of choosing the test cards so that a change in just one attribute is made after the subject happens onto a positive instance. For example, suppose the subject chooses a card with two green circles and two borders. He might then decide to concentrate on the attribute of shape, and for his next card choose another shape to test (for example, two green crosses with two borders). Or finally, the subject might take a gamble and change two or more attributes at the same time. This the investigators call *focus gambling*; it is gambling because, while the subject may solve the problem in fewer trials than with conservative focusing, he may also take longer.

The use of strategies. In general, as we might suppose, subjects do better with conservative focusing than with logical strategy (or an approach to it). The superiority of the focusing strategy becomes greater when the strain placed on memory and assimilation becomes greater. People who try to approximate the logical strategy do more poorly as conditions become more difficult. If the subjects are limited in the number of choices they are permitted to make (imposed by telling the subject that he can have only four or perhaps even one choice to try to solve the problem), the subjects tend to shift over to focus gambling. Finally, subjects change their strategy as they move from problem to problem. Subjects who have had many positive examples in earlier problems gradually shift to a focus gambling strategy, while those subjects who have had many negative examples shift more toward the conservative, one-change strategy. Thus, concept attainment strategies are altered by the subject's subjective estimates of changes in the probability of payoff.

Thus, even such an extended and general characteristic of behavior as the use of strategy in problem solving and concept attainment is determined by the pattern of previous reinforcement. Also individual habits, less directed by payoff than by the verbal mediational activity characteristic to one's life and work, influence strategy. Bruner, Goodnow, and Austin point out that a few of their subjects of very superior intelligence and usually with some mathematical training tried to attain the concepts by a purely logical strategy. This was an error in the sense that it failed to meet the objectives of a workable strategy. Thus the mediational processes that directed these

³ An experiment by Whitfield (1951) shows that the same limitations of assimilation and memory prevent human beings from being perfectly efficient in the solution of trial-and-error problems.

people to attempt a purely logical strategy failed to produce adaptive behavior.

Learning and Thinking

In conclusion, we can point out that the intimate relationship between learning and thinking or problem solving does not mean that efficient thinking is entirely determined by associative mechanisms. It is, however, certainly determined to a considerable extent by associative learning, and furthermore the limitations of associative learning provide very real and important limits to our ability to solve problems. We cannot behave like logical computing machines, because we do not have the ability to scan our memories almost instantaneously and pick out and act upon the appropriate stored information. We can carry in our heads a limited amount of information, and we are dependent upon external cues and the situations of the moment for our ability to extract and act upon any piece of information in our memory. Furthermore, from the logical point of view our storage systems are haphazard and inefficient. We recode things in a hit-or-miss fashion; we may rely on a nonsense rhyme, for instance, to help us remember the order of the cranial nerves for an examination in comparative anatomy.

We trade order and magnitude for flexibility, however. A computing machine can operate on only one set of instructions. It is saddled with the strategy set into it by its operators. Its human operators, on the other hand, can move here and there. They can make guesses and govern the reasonableness of their guesses by an intuitive and inefficient, but effective, calculation of subjective probabilities of payoff. These guesses and the strategies for dealing with them are learned, and the ability of human beings to make effective guesses in working out problems is determined by the pattern of past associative learning.

At the same time we are at the mercy of our prejudices and beliefs. A large experimental literature is devoted to the demonstration that critical thinking and logical reasoning among individuals often turns out to be a reflection of prejudice and earnestness of belief.

These dynamic factors, as we have seen, are an important part of associative learning. Motivation influences behavior, and it is behavior, overt or implicit, that we learn. Therefore, we learn what our motives dictate that we learn. In part, this means that our thinking processes are indirectly determined by the influence of dynamic factors on learning, but it also means that these processes are directly determined by the same factors that determine learning and memory.

CHAPTER 12

THEORY AND APPLICATIONS OF LEARNING

Although we have examined many specific theoretical problems in the psychology of learning, we have not yet taken the opportunity to look at some of the more general theoretical questions. The purpose of this chapter is to survey some of these broader aspects of learning.

In the course of examining particular topics such as partial reinforcement, we have seen some of the hypotheses that have come out of various theories of learning. In this chapter, rather than concentrate on these limited hypotheses, we shall look at some of the contrasting strategies of theories of learning and how these came to be. In addition, we shall want to consider the relationships of the study of learning to other biological and social disciplines. For example, the neurophysiological basis of learning is one of the most fundamental problems in physiological psychology, and we could not consider our survey of learning complete without at least a glance at current neurophysiological theorizing. Another important question concerns the relations between animal behavior in nature and the facts and theories of learning. We should be neglecting something fundamental if we did not briefly survey some of the current work on this topic. On the other side of psychology, we have yet to consider the psychology of learning in relation to personality theory, social behavior, and the educative process. These are topics we shall need to cover.

This chapter, then, will include a summary of the theoretical background of the psychology of learning and of the relationships between the study of learning and other biological and social disciplines. In the first section we

shall consider learning theory and after that, turn to a survey of the broader implications of the psychology of learning.

LEARNING THEORY

Psychology generally, and particularly the psychology of learning, has long been noted for the exotic variety of its theories and the bitter arguments of their exponents. For several reasons, however, we shall not examine each of the classical and current theories of learning. First of all, this task has already been done in a very complete way by Hilgard (1956), and secondly, the grand, universal theoretical controversies are somewhat out of date. Not too many years ago, a psychologist interested in learning was identified by his adherence to one current theory or another, and these theories were "opposed" to one another from the ground up. Theoreticians (and experimentalists) spent much time in fruitless argument with one another, either by damning the logic of an opposing theory or by attempting to demolish it through the performance of "crucial" experiments.

Such theoretical argument has not entirely disappeared today for the good reason that hardly any of the basic issues have been settled to everyone's satisfaction. The strong emphasis upon large scale theories of behavior and devoted adherence to one theory or another have disappeared. Currently, psychologists in learning are more likely to quarrel about methods of approach and about areas fruitful for research than about universal theories. Since it is recognized that there is no theory of learning currently available which is able to "explain" all the facts of learning, there is little talk about the exclusive and ultimate correctness of one or another theory.

One of the things likely to arouse an argument among psychologists is the question of the relative merits of trying to derive from rigorously quantitative deductions limited hypotheses about highly specialized problems or trying to sketch in a rather intuitive fashion a qualitative description of fundamental problems and those of general interest. Thus, one group of psychologists may be interested in rigorously deriving a mathematical theory of the effects of partial reinforcement on the behavior of the rat in the runway, and another group may be interested in trying to find the main elements in the relationship between associative learning and problem solving in intelligent human adults. There is really little to argue about between such groups, except the relative merits of the problems they have chosen to work on, since the problems themselves will only at some future date come to a common ground.

In this section we shall examine a few of the approaches that are currently important in the theory of learning. This will not be an exhaustive account, since we have space only to discuss a limited number of problems and methods that are of interest to theorists at present.

The Quantitative Approach

The Hullian School

In the background of the main currents of attempts to derive theoretical principles of learning by mathematical techniques stands the work of the late C. L. Hull. His chief efforts in the theory of learning are only in a very primitive sense mathematical, but he, more than any other worker, alerted psychologists to the possibilities of applying mathematics to quantitative problems in learning.

Hull's approach. Hull's method¹ was to make a number of postulates and see what theorems could be derived from combinations of these. Postulates are fundamental statements not directly testable by experimental means. Theorems, which are derived by rigorous logic from the postulates, are experimentally testable, however. Thus, one of Hull's postulates concerns a stimulus trace, or afterdischarge, as the result of stimulation. Such a trace cannot be directly measured. Its properties in conjunction with other postulated properties lead to predictions about the form of the relationship between ease of conditioning and time intervals between stimuli. This relationship should be testable.

Though this approach provided the basic framework for Hull's work, in actual practice his theorizing fell considerably short of his aims (Koch, 1954). Very often, he slipped into attempts to demonstrate the adequacy of his postulates by citing experimental evidence to show how "reasonable" they were. His deductions were sometimes contradictory and in logical error. Furthermore, his theory seemed more quantitative than it really was. This statement deserves some further comment.

Hull's quantitative theorizing consisted mostly of putting his basic postulates in the form of mathematical equations. He did little in the way of deriving theorems by mathematical means from these postulates. Ordinarily it is the mathematical derivation of theorems from postulates that people mean when they talk about mathematical theories. Hull simply states an equation and lets it go at that. Since there are almost limitless possibilities

¹ The best statement of Hull's beliefs about the method and philosophy of science is in his *Principles of Behavior* (1943).

for the choice of equations for each situation, his efforts in this direction appear arbitrary to say the least.

An even greater difficulty with Hull's quantitative theorizing is the fact that his basic equations seem to lead nowhere. Often, psychologists state relationships in a quasi-mathematical form. Thus we may have heard "tension is directly proportional to the amount of conflict." Such a statement implies that the relationship between tension and conflict can be written

$$T = kC$$

Now there is nothing wrong with this statement and the equation except that we have no way to measure either tension or conflict in such a way that we can confidently test an apparent linear relationship between them. There is just no way to measure quantitatively either amount of tension or amount of conflict.

This inability to measure many of the concepts Hull talked about was part of his difficulty. His most important hypothesis about learning is that of *habit strength*, which is an intervening variable not directly measurable. In order to make a quantitative statement about habit strength that can be tested experimentally, it must be related to direct measures of behavior such as the amplitude of a galvanic skin response. Since there are an enormous number of measures of behavior that can be used to evaluate learning, there must be an equally large number of arbitrary "postulates" needed to relate habit strength to behavior. Hull made valiant efforts to solve this problem (Hull, et al., 1947), but his attempts were so conspicuously unsuccessful that they reflected on the adequacy of his whole theory.

Hull's method provided two important lessons for subsequent quantitative theorists. (1) Given the present confusion over basic problems of measurement and fact in psychology, a rigorous postulational approach can work only by making each theory cover a very limited area. A comprehensive deductive theory of behavior is practically impossible at present. (2) The mathematical relationships in the theory must be developed from a very few definitions and postulates. The proliferation of arbitrary mathematical equations simply leads to confusion, both in the theory and in the experimental test of the theory.

Hull's contribution. Although we have been critical of Hull's method, it should not be assumed that all his work was futile. In the first place, he *did* introduce learning psychologists to quantitative theory. Secondly, and more importantly, Hull fixed for many years the basic problems and concepts of learning theory. Hull made a fruitful union of mathematics, classical association theory (see page 309), and motivation theory—and out of this

came our basic contemporary concepts of drive, stimulus generalization, secondary drive, etc. Thus, to a considerable extent we owe the basic conceptual framework of the psychology of learning to Hull. Workers in motivation theory, in the physiology of motivation and learning, and in personality theory—all use concepts which were either introduced or made popular by him.

Not the least of Hull's contributions was in forcing other theorists to come to grips with him. The long-standing controversy between Hull and his associates and Tolman and his associates led up many a blind alley, but out of it came important knowledge. This controversy contributed to the clearer understanding of the separation of motivational and learning variables and to the understanding of the role of reinforcement in controlling behavior. Therefore, in a very real sense, Hull stands at the origin of much of modern behavior theory.

Work in the tradition of Hull is very much alive. K. W. Spence has attempted to correct many basic defects in Hull's theory. He has severely limited the applications of his fundamental postulates to very special problems in behavior. He has, at the same time, attempted to simplify Hull's basic theory and make it more amenable to ordinary mathematical treatment. In a recent book (Spence, 1956) he has presented his own theory of learning, and this probably stands as the final and culminating work of the Hullian school.

Statistical Learning Theory

The behavior of animals and men is variable. Learning curves are irregular. Most earlier learning theorists regarded this variability as a nuisance which could be minimized by careful experimental control and by studying behavior only in the most circumscribed situations. There exists, however, a mathematics of variability, or randomness, and it was inevitable that sooner or later psychologists would apply some of the mathematics of randomness to the study of behavior. Hence, the statistical theories of learning.

The basic notion behind any statistical theory is that the events one measures in the laboratory are the result of the influence of a whole population of elementary random events. These are events which cannot be predicted; a good example is the unbiased toss of a coin. If we toss a penny enough times it will come out heads about 50 per cent of the time and tails about 50 per cent of the time. We can think of the percentage of heads as the "result" measured in the laboratory and each individual toss of the coin as one of the elementary events that goes to make up the result. In statistical

learning theory, postulates are stated about the distribution of elementary random events, and measurable results are predicted from a mathematical development of these postulates.

In recent years a number of theorists have developed statistical learning theories. The theories of Bush and Mosteller (1955), Estes and Burke (1953), and Restle (1955) are all rather similar to one another in method and outcome. Estes and Burke (1953) have called theirs a theory of "stimulus variability." The essential notion they use to derive the basic curves for classical and instrumental conditioning is that each time a response occurs, it becomes associated with a *sample* of all of the possible stimuli that are likely to impinge on the organism on successive trials. We can think of the sample as a random sample of stimulus elements that manage to get through the sensory input of the organism at the time a particular response occurs.

Notice that in this theory, learning is an all-or-nothing affair. The sample of stimulus elements that occurs simultaneously with any given response becomes conditioned to that response. The reason, according to the theory, that learning curves in practice are gradual and irregular is that different samples occur on successive trials. These samples will overlap to some extent, so after an initial trial some of the stimulus elements will already be conditioned. Others will not, however, and these become conditioned on that trial. We can see that if the population of stimulus elements were a limited one, eventually so many elements would have been conditioned that there would be none left to condition. Thus, conditioning approaches a limit of 100 per cent. An intuitive idea of how the cumulation of learning occurs can be derived from Figure 57.

The strategy of statistical learning theories. Statistical learning theorists have, to a considerable extent, avoided the problems which plagued Hull. Their postulates are for the most part elegantly simple. Bush and Mosteller (1955), for example, are able to state an equation which implies that the change in probability of a response will be determined by two parameters, one associated with reward and the other with punishment. The rest of the mathematical development of their theory is given over to the applications of their basic equation and the definition of events, such as trials, and measurements, such as latency.

Furthermore, statistical learning theorists are not bothered by the cumbersome necessity of transforming equations which relate independent variables to intervening variables into equations which relate independent variables to dependent variables. The equations are all stated in such a form that

there is only one fundamental measure of behavior, the probability of occurrence of a particular response. All other measures, such as magnitude of response and latency of response, are derived from definitions of probability of response. Therefore, the statistical learning theories have the considerable merit of suggesting a fundamental measurement in psychology, the measurement of probability or, empirically, the relative frequency of occurrence. This has found widespread acceptance among theoretical psychologists, for a variety of efforts in psychology have recently converged on the common notion that the basic measurement of behavior must be in terms of probability of occurrence of different responses.

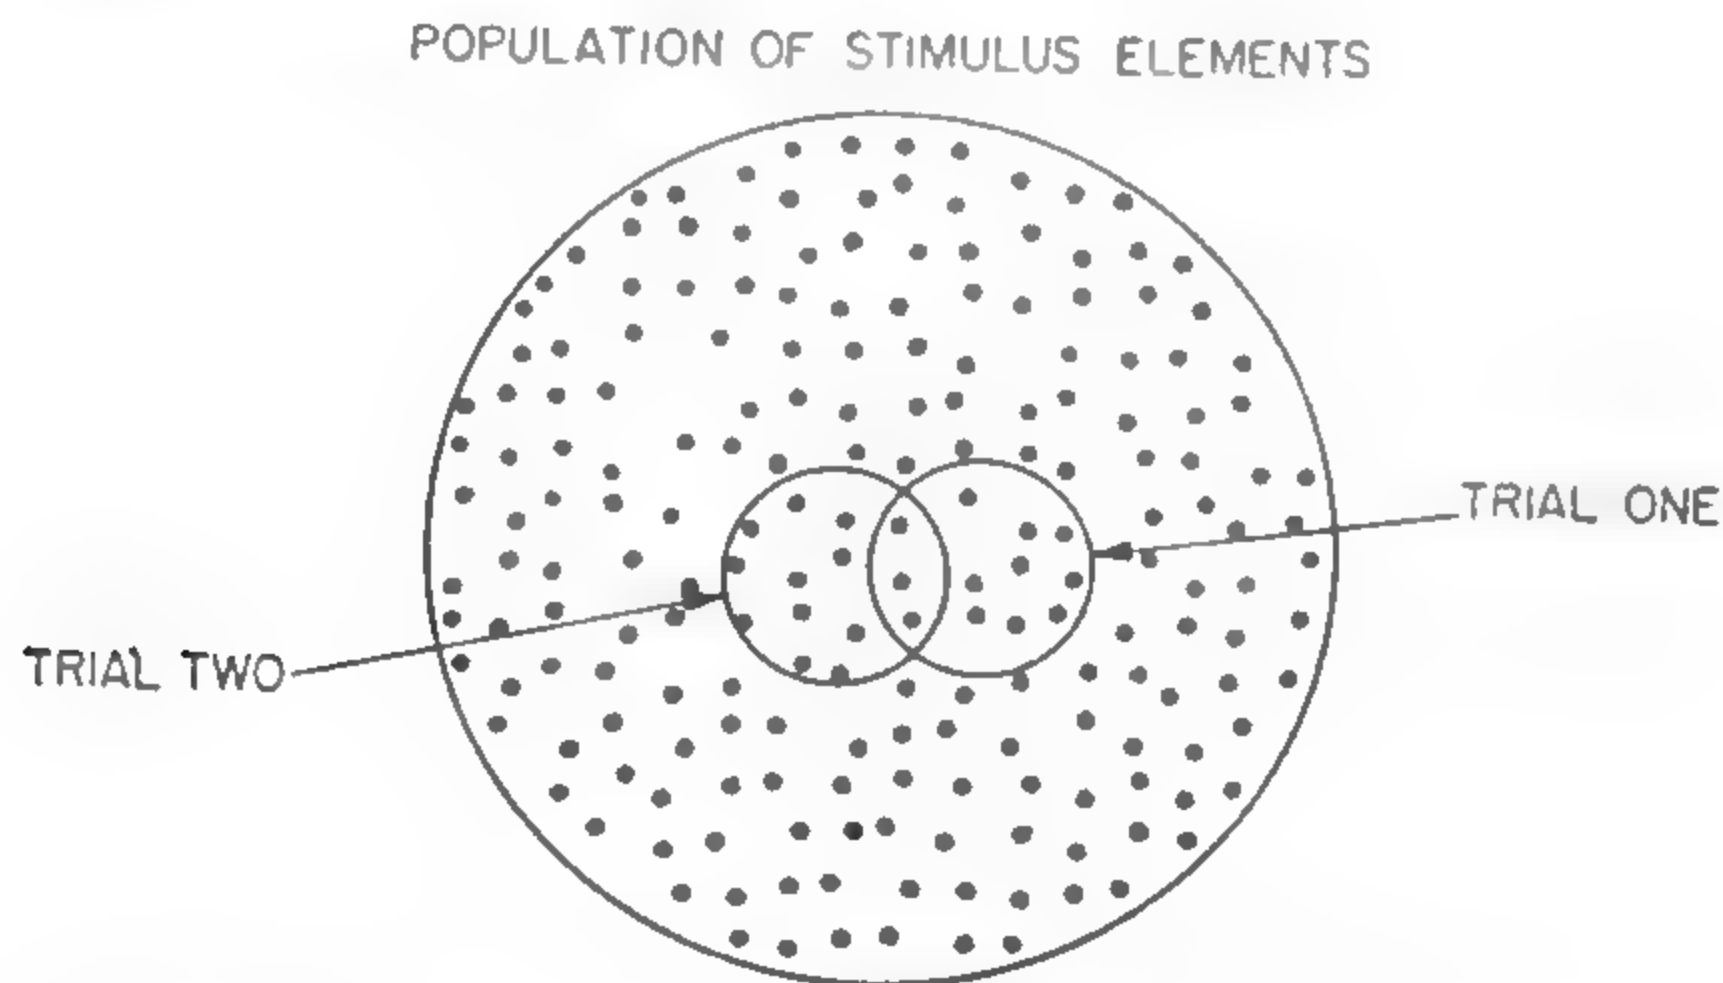


Figure 57. An illustration of two successive trials according to statistical learning theory. Each dot represents one of the "elements" of the learned act. The dots enclosed by trial 1 are present and hence conditioned on the first trial. The dots enclosed by trial 2 are present on the second trial. Notice that some of these dots are already conditioned, so that the second trial adds fewer elements to the strength of conditioning.

The limitations of statistical learning theory. The elegance and sophistication of statistical learning theories should not blind us to certain defects. One concerns the kinds of behavior the proponents of these theories have chosen to study. To a considerable extent, these theorists have gone over the same ground covered earlier by the Hullian school. This has meant that they have not discovered much that is exciting or new.

Strangely enough, where these theorists have ventured into new fields, there has been considerable dissatisfaction with the results. One field that has been invaded to quite an extent by the statistical learning theorists is the study of "guessing behavior" in human beings. The experiments on such behavior are familiar to us, since we reviewed them in the section on partial reinforcement. It is the fact that these problems grew up as analogies to the

effects of partial reinforcement in animal behavior that has led to the attempt to deal with them by statistical learning theories.

In the guessing experiments, subjects face two lights mounted on a board. The subjects are to guess while one light is on whether the other one will come on also. The experiment is arranged so that this will occur a certain number of times, though the individual occurrences are nearly random. In this experiment, subjects are likely to guess, after much practice, about the same percentage as the actual occurrences of the light. This effect can be derived from statistical learning theory. Unfortunately, however, this same effect can be derived by a number of other more complicated, but perhaps more reasonable, theories. Furthermore, the effect does not always occur. It is not the rational strategy, and statistical learning theory has a most unconvincing way of handling the rational strategy.

Some experiments have been designed specifically to test the implications of statistical learning theory for guessing behavior (Estes and Straughan, 1954). These experiments are very limited in scope, however, and the test of the theory consists of curve fitting based on empirical constants, a technique rather more likely to lead to success than failure.

These guessing-game experiments are likely to strike the outsider as rather sterile and without significance to the basic problems of human behavior. As models of human decision making they have a certain merit, but the statistical learning theories so patently ignore some basic facts of human behavior (the ability of human beings to remember several guesses previous to the one they have just made, for example) that the use of these experiments as tests of the theory seem to lead nowhere. The theory and the experiments are, therefore, much more convincingly applied to the behavior of rats.

A somewhat different approach gives promise of some merit. By using many of the same notions as those used by the statistical learning theorists, Miller and McGill (1952) have been able to construct a small theoretical model of free recall of verbal items. This theory assumes that the probability of recalling a word on any trial is completely determined by the number of times the word has been recalled on previous trials. Such an assumption appears to be quite reasonable as a first approximation to a quantitative theory of verbal recall. Miller and McGill examine the implications of several levels of complexity of the theoretical model, and again by curve fitting, the theory seems to fit some data (it should be mentioned, however, that in this case the theory leads to a rather novel way of treating and presenting data). Unfortunately, this theory must introduce some assumptions that

are patently contrary to fact. For example, it must assume that the different words used all have independent probability of being recalled.

In general, both the weakness and strength of statistical learning theory lies in the introduction of oversimplification. It is a strength because it suggests fundamental variables which need clear and explicit treatment before we can go on to the more complicated problems, and it is a weakness because the behavior predicted by the theory is limited to special and exotic cases only remotely related to the behavior of real men and animals.

Two General Approaches to the Psychology of Learning

Association Theory

If the human being is the featherless biped, he is even more characteristically the verbal animal. The major effort in trying to understand the nature of learning in man has been directed toward verbal learning. Experimental study of verbal learning has been enormously influenced by the tradition of associationism, and in order to understand some of the current pre-occupations of psychologists interested in verbal learning, we need to look briefly at classical associationism and how it has influenced current work.

Actually, associationism provides the deepest current in the psychology of learning, and nearly all theorizing has either been directly in the associationistic tradition or influenced by it. Association theory has its beginnings in Aristotle, and it runs very strongly through the history of British philosophy from the seventeenth to the nineteenth century. The doctrine of the "association of ideas" was designed to provide a description of the continuity of mental life. Ideas provide the units, or elements, of mental life, and they follow one another in a constant chain. One particular idea follows another one because, sometime in the past, these ideas have occurred together. Or one idea elicits another one because they are similar or contrasting.

Classical association psychology has disappeared, but the idea of the discrete association of elements of behavior by contiguity and by similarity (stimulus generalization) provides some of the fundamental elements of the modern psychology of learning.

Relationships between learning, motivation, and incentive conditions in modern psychology have much of their origin in the amalgamation of associationism and hedonism put forward by the British philosophers. Locke, for example, reminds us that the ideas that make the deepest and most lasting impression are those accompanied by pleasure and the relief of pain. Thus

he says, "the great business of the senses being to make us take notice of what hurts or advantages the body, it is wisely ordered by nature that pain should accompany the reception of several ideas." ²

We have seen in earlier chapters that Ebbinghaus successfully translated the doctrine of association of ideas into experimental work on the higher mental processes. Since Ebbinghaus's time there has been a continuous effort to expand and elaborate the experimental analysis of the associative processes. From Robinson (1932) to McGeoch (1942) to Underwood (1957), the experimental work on association has contributed to the scientific understanding of human nature.

Why association theory? Why has so much attention been paid to association theory by psychologists interested in verbal learning? Perhaps the most important reasons why associationism is an ancestor of so much of the contemporary psychology of learning while other philosophic doctrines have disappeared are (1) its emphasis upon cause and effect and the chaining of ideas and (2) its analytic nature. Since ideas and the words that accompany them can be separated one from another and because they can be combined in every imaginable way, it is possible to treat arrangements of them as independent variables. Independent changing of the controlling variables is the cornerstone of the experimental method, and classical association theory provided a framework for the experimental investigation of verbal behavior.

Thus the hallmark of associationism in the modern study of verbal learning is the independent variation of the stimuli or words presented to the learner. We put nonsense syllables, etc., together in different ways and see what effect the different arrangements have upon the behavior of the learner. In so doing, we find justification for emphasis upon contiguity, repetition (the "law" of association added by Hartley), and similarity in determining how behavior will organize itself. With the additional consideration of factors of motivation and incentive, a complete program for the experimental investigation of verbal learning emerges.

The limitations of the associationistic approach. The classical associationistic approach to the study of verbal learning and memory is exemplified by the methods of rote serial learning. These methods are designed to exert the greatest control over the learner; they force him to associate items according to the design of the experimenter, not according to his own wishes. The result is that the typical experiment on the verbal associative

² John Locke, "Concerning Human Understanding," in E. A. Burtt, *The English Philosophers from Bacon to Mill*, Modern Library, Inc., New York, 1939.

immediate past, he assigns an important role to neurophysiological hypotheses and makes them the central element in his theory. As part of a general neurophysiological theory of behavior, Hebb has suggested a theory of learning, and from this we can abstract some of the important ideas concerned with changes in the nervous system. Thus we can see how one psychologist has imagined events to take place in the individual nerve cells for during learning.

Hebb begins by pointing out that we must have some way of reconciling the seemingly incompatible demands of perception and learning. Learning apparently requires that *specific* cells be excited for a specific response to be remembered. Perception, on the other hand, seems to demand that patterns of excitation, rather than specific cells, be the physiological basis for psychological events. We know, for example, that a square looks like a

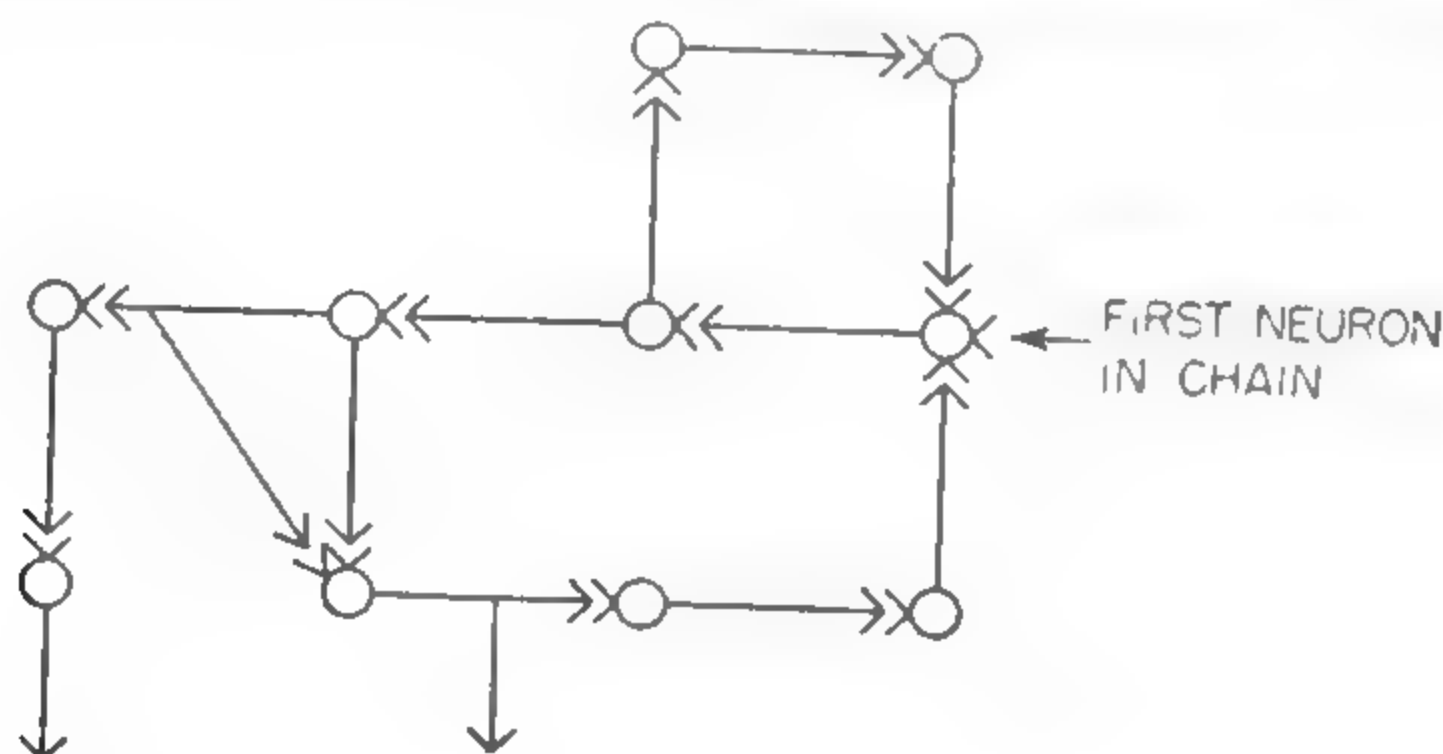


Figure 58. A schematic diagram of a recurrent nerve circuit.

square whether its image on the retina forms a square or a diamond. Furthermore, the same square is always immediately recognized as the same square, even if it excites different portions of the retina, and hence different portions of the central visual system.

Hebb solves this dilemma by making learning basic to perception. He suggests, for example, that we must *learn* to perceive the square as a square no matter how it is viewed. The problem then remains to postulate the manner in which specific cells become the basis for memory. Hebb does this by suggesting a dual-trace mechanism—two factors in the development of the memory for specific responses. One factor is built on the principle of recurrent nerve circuits (schematically illustrated in Figure 58). We can see that excitation started by the first nerve cell in the chain returns to that cell through a number of other cells. We know that recurrent circuits are active for brief periods of time following stimulation of the central nervous system. A recurrent circuit can act to preserve for a brief period of time excitation in a particular group of nerve cells. Thus, it can temporarily

hold the memory of some stimulation until a secondary but permanent connection has developed sufficiently to take over. The secondary change (the one responsible, says Hebb, for everything but events such as the immediate memory span) is the result of growth of nerve fibers toward one another at the synapse, or junction, between the nerve cells. He specifically suggests that one cell acts upon another to form knoblike growths that make the transmission of impulses through one particular group of cells easier than through another. The growth of the fibers itself is the result of activity of the nerve cells. Thus, when a particular conjunction of external stimuli occurs, they are likely to elicit activity in the same group of cells that was previously active. The result is a permanent connection between the external stimuli and events in the central nervous system.

The Potentiation Effect and Eccles's Theory

This second, permanent mechanism for establishing connections between cells is very much like a mechanism suggested by the neurophysiologist Eccles (1953). Since Eccles derives his notions from direct neurophysiological evidence, it would appear that we are remarkably close to closing the gap between neurophysiology and the psychology of learning.

Facilitation of one reflex function by the activation of another is a common neurophysiological observation. Ordinarily, however, such facilitation is limited to the time at which the activation of the second system occurs or to a fraction of a second thereafter. Recently, however, neurophysiologists (Lloyd, 1950; Eccles and Rall, 1951) have observed a facilitation effect which lasts not for a few thousandths of a second but for a good many minutes. This effect has been called *post-tetanic potentiation*. On a logarithmic time scale this long-lasting facilitation is well on the way to the relatively permanent facilitation of responses to conditioned stimuli produced by their being paired with unconditioned stimuli.

Eccles believes the potentiation effect to occur by a swelling of the end knobs at synaptic junctions (this is, of course, similar to Hebb's notion of end-knob growth). Eccles then goes on to develop a theory of conditioning built on the mechanism (end-knob swelling) that seems to account for post-tetanic potentiation.

His theory of conditioning concerns the interaction between systems of neurons called *pathways*. These pathways converge on one another in such a way that the action of one tends to facilitate action in another. Thus the occurrence of excitation in the pathway carrying impulses from the unconditioned stimulus facilitates the impulses from the conditioned stimulus so

that they carry over to the motor neurons and produce a conditioned response.

The theory can account for some of the facts of classical conditioning (the fact, for example, that the conditioned response is not quite the same as the unconditioned response), but it is wide of the mark on many other facts. For one thing, the theory clearly predicts that a conditioned response should be weakened through disuse, and that, indeed, disuse should be a more effective weakening agent than experimental extinction. This is clearly contrary to fact, and presents a serious theoretical difficulty. Malmö (1954) has suggested that this could be corrected by appealing to a neurophysiological principle of inhibition (hyperpolarization). Thus it is quite likely that additional work may make it possible for the theory to account more adequately for the basic facts of conditioning.

Although the theory does not adequately explain the behavioral facts, it is solidly based on neurophysiological concepts, and here lies its importance. It is not an *ad hoc* theory put together to account for the behavioral facts of conditioning; it is a theory that attempts a genuine integration of neurophysiological concepts and the principles of conditioning. It is possible that this represents the first real bridge between behavior and the basic laws governing the interaction between neurons in the central nervous system. At any rate it is difficult to see how the long-lasting facilitation effects found in post-tetanic potentiation can be ignored in any future neurophysiological theory of conditioning.

Ethology and the Psychology of Learning

Experimental psychologists, for the most part, have been interested in the study of the basic principles of behavior, particularly as they apply to the behavior of human beings. Biologists and naturalists, on the other hand, have been interested in the behavior of animals and in the relation between animal behavior and biological problems in ecology and evolution. This kind of study has come to be known as *ethology*. Quite recently, the laboratory approach to behavior, characteristic of the experimental psychologists, and the field approach, characteristic of the ethologists, have tended to come together. In this section we shall deal with some of the relations between ethology and the psychology of learning.

Habituation and Imprinting

We have described the two basic kinds of learning to be classical conditioning and instrumental learning. We have, however, restricted our dis-

cussion almost entirely to the behavior of mammals. When we consider learning in other vertebrates and in invertebrates, the ethologist reminds us that there are kinds of learning which do not seem to be easily described as either classical or instrumental conditioning. Two other classes of the modification of behavior that have been intensively studied by students of animal behavior are habituation and imprinting, and we shall consider these briefly.

Habituation. Humphrey (1933) gave us the first clear account of habituation and its importance in the animal world. Habituation is learning not to respond to stimuli which tend to be without significance in the life of the animal (Thorpe, 1956). A familiar example is the description of a simple experiment with a snail (Humphrey, 1933). If as the snail moves along a board, the board is periodically jerked, there is a reflex withdrawal of the tentacles. As the jerking is repeated, however, this withdrawal movement gradually diminishes and finally stops altogether. If the intensity of the stimulus is increased, the response will reappear, only to disappear with further repetitions. A period of rest will allow the response to recover spontaneously, but as in the case of the spontaneous recovery of extinction, repeated experimental sessions will finally cause the behavior to disappear more or less permanently. This last point is especially important, because it is this which differentiates habituation from simple fatigue.

It is extremely difficult to obtain evidence for associative learning (either of the classical conditioning or the instrumental variety) in single-celled animals or in very primitive metazoa, such as the coelenterates. On the other hand, habituation can be readily demonstrated with every animal (Thorpe, 1956) and it must be regarded as the most fundamental and elementary example of the permanent modification of behavior by repeated stimulation. Of course, habituation is very similar to experimental extinction, and it may well be that the most elementary examples of extinction are examples of habituation of learned rather than unlearned responses.

Thorpe (1956) points out that habituation is of fundamental importance in the natural life of animals. Furthermore, he emphasizes, it is probably related to the gradual atrophy of instinctive reactions when they are allowed to happen without consummation of the appropriate appetite behind them. Thus, both learned and unlearned acts seem to be subject to disappearance on repeated elicitation. It is only the activation of the motivation system (perhaps centered in the hypothalamus) that keeps responses alive. In other words, widespread habituation of behavior may very well be something which has to be overcome by the positive effects of motivation.

Imprinting. In recent years there has been a good deal of attention given to a peculiar variety of associative learning called *imprinting*. It was originally discovered in birds, but it is now suspected to be quite widespread in the animal kingdom. The peculiar features of imprinting are as follows: (1) it occurs with spontaneous rather than reflex behavior (though it is behavior "released" by external stimulation), (2) it requires no specific reinforcement for activation, and (3) it seems to occur only during a restricted period in the life of very young animals.

The phenomenon of imprinting first received attention from the work of Lorenz (1935), who described a typical observation on imprinting. The eggs of a goose are hatched in an incubator and the young are restricted to the influence of human beings only; they never see the adults of their own species. The result is that the goslings come to follow the human being (when he quacks), rather than one of their own species. As we know, goslings typically follow their mother in a procession. This particular activity is not so much the result of the distinctive visual stimuli from the parent bird as it is the result of the association of any moving object the bird sees with the release of the response of following. This imprinting of the young bird appears to be limited to very brief periods in its life (Ramsey and Hess, 1954).

It has sometimes been supposed that imprinting is irreversible both in the sense that it is never forgotten and in the sense that once a young bird is imprinted to respond to a particular stimulus (say a moving shoe box) it cannot be imprinted to other objects. The latter statement is apparently untrue; Jaynes (1956) was able to obtain considerable evidence of generalization of imprinting to objects quite dissimilar to the original imprinted stimulus in young chicks.

Imprinting is the acquisition of the ability by new cues to command innate behavior. Although most of the evidence on imprinting has been obtained with birds, there is evidence that a similar effect occurs in mammals and perhaps in insects (Thorpe, 1956). The great importance of imprinting in the life of animals is the plasticity, or freedom, in response to stimuli that it gives to certain kinds of instinctive activities. We are still uncertain about how it is related to the more universal kind of associative learning exemplified by classical conditioning and instrumental behavior. Until we know more about it, imprinting remains something of a curiosity in the description of learning.

Ethology and Learning

Thorpe (1956) in his comprehensive survey of the relationships between the mechanisms of learning and the natural behavior of animals shows how particular aspects of learning such as stimulus generalization, latent learning, trial and error, etc., fit into the biological pattern of survival in the natural world. He demonstrates how the mechanisms of latent learning work in exploratory activity and in the learning of topography in animals as primitive as insects.

He describes, for example, the dependence upon learning of an instinctive bit of behavior, the egg-laying and provisioning behavior of the hunter wasp. This animal is able to make several burrows in different places, lay eggs in each of them, and return later to provide food for the young. The tendency to return is itself determined by innate releasing mechanisms (instinct), but the location of the burrows must be learned by the animal. The learning is latent because the behavior of returning to the burrow is not always accompanied by direct reinforcement.

Furthermore, Thorpe was able to demonstrate, by controlled observations, that this must be the learning of stimulus relations rather than specific responses, for these animals can be picked up and released in strange locations and still find their way back to the nest (Thorpe, 1950). Figure 59 shows some of Thorpe's tests for detour, or barrier-circumventing, behavior in these wasps as well as the effects of release in a strange location. Notice the straight and regular path of the wasp to the burrow following its release in new surroundings.

Thorpe (1956) argues that some of the spectacular observations on communication in bees that have received so much attention recently must indicate rather complicated latent learning in these animals. Bees can return to a hive after locating a source of sugar, and when they return they communicate to other bees the direction and approximate distance of the source. The communication is by means of a complicated dance. Dethier (1957) has shown that a similar pattern of behavior in the fly can be accounted for by direct stimulus-response and physiological mechanisms, and it is very likely that the activity of the bee can be accounted for by similar mechanisms. Thus, spectacular as this behavior is, it is probably not an example of learned behavior.

There is great difficulty in separating the influence of instinct and learning in animals. Only now are we beginning to understand how these things interact in the natural world. Most animal behavior is a complicated inter-

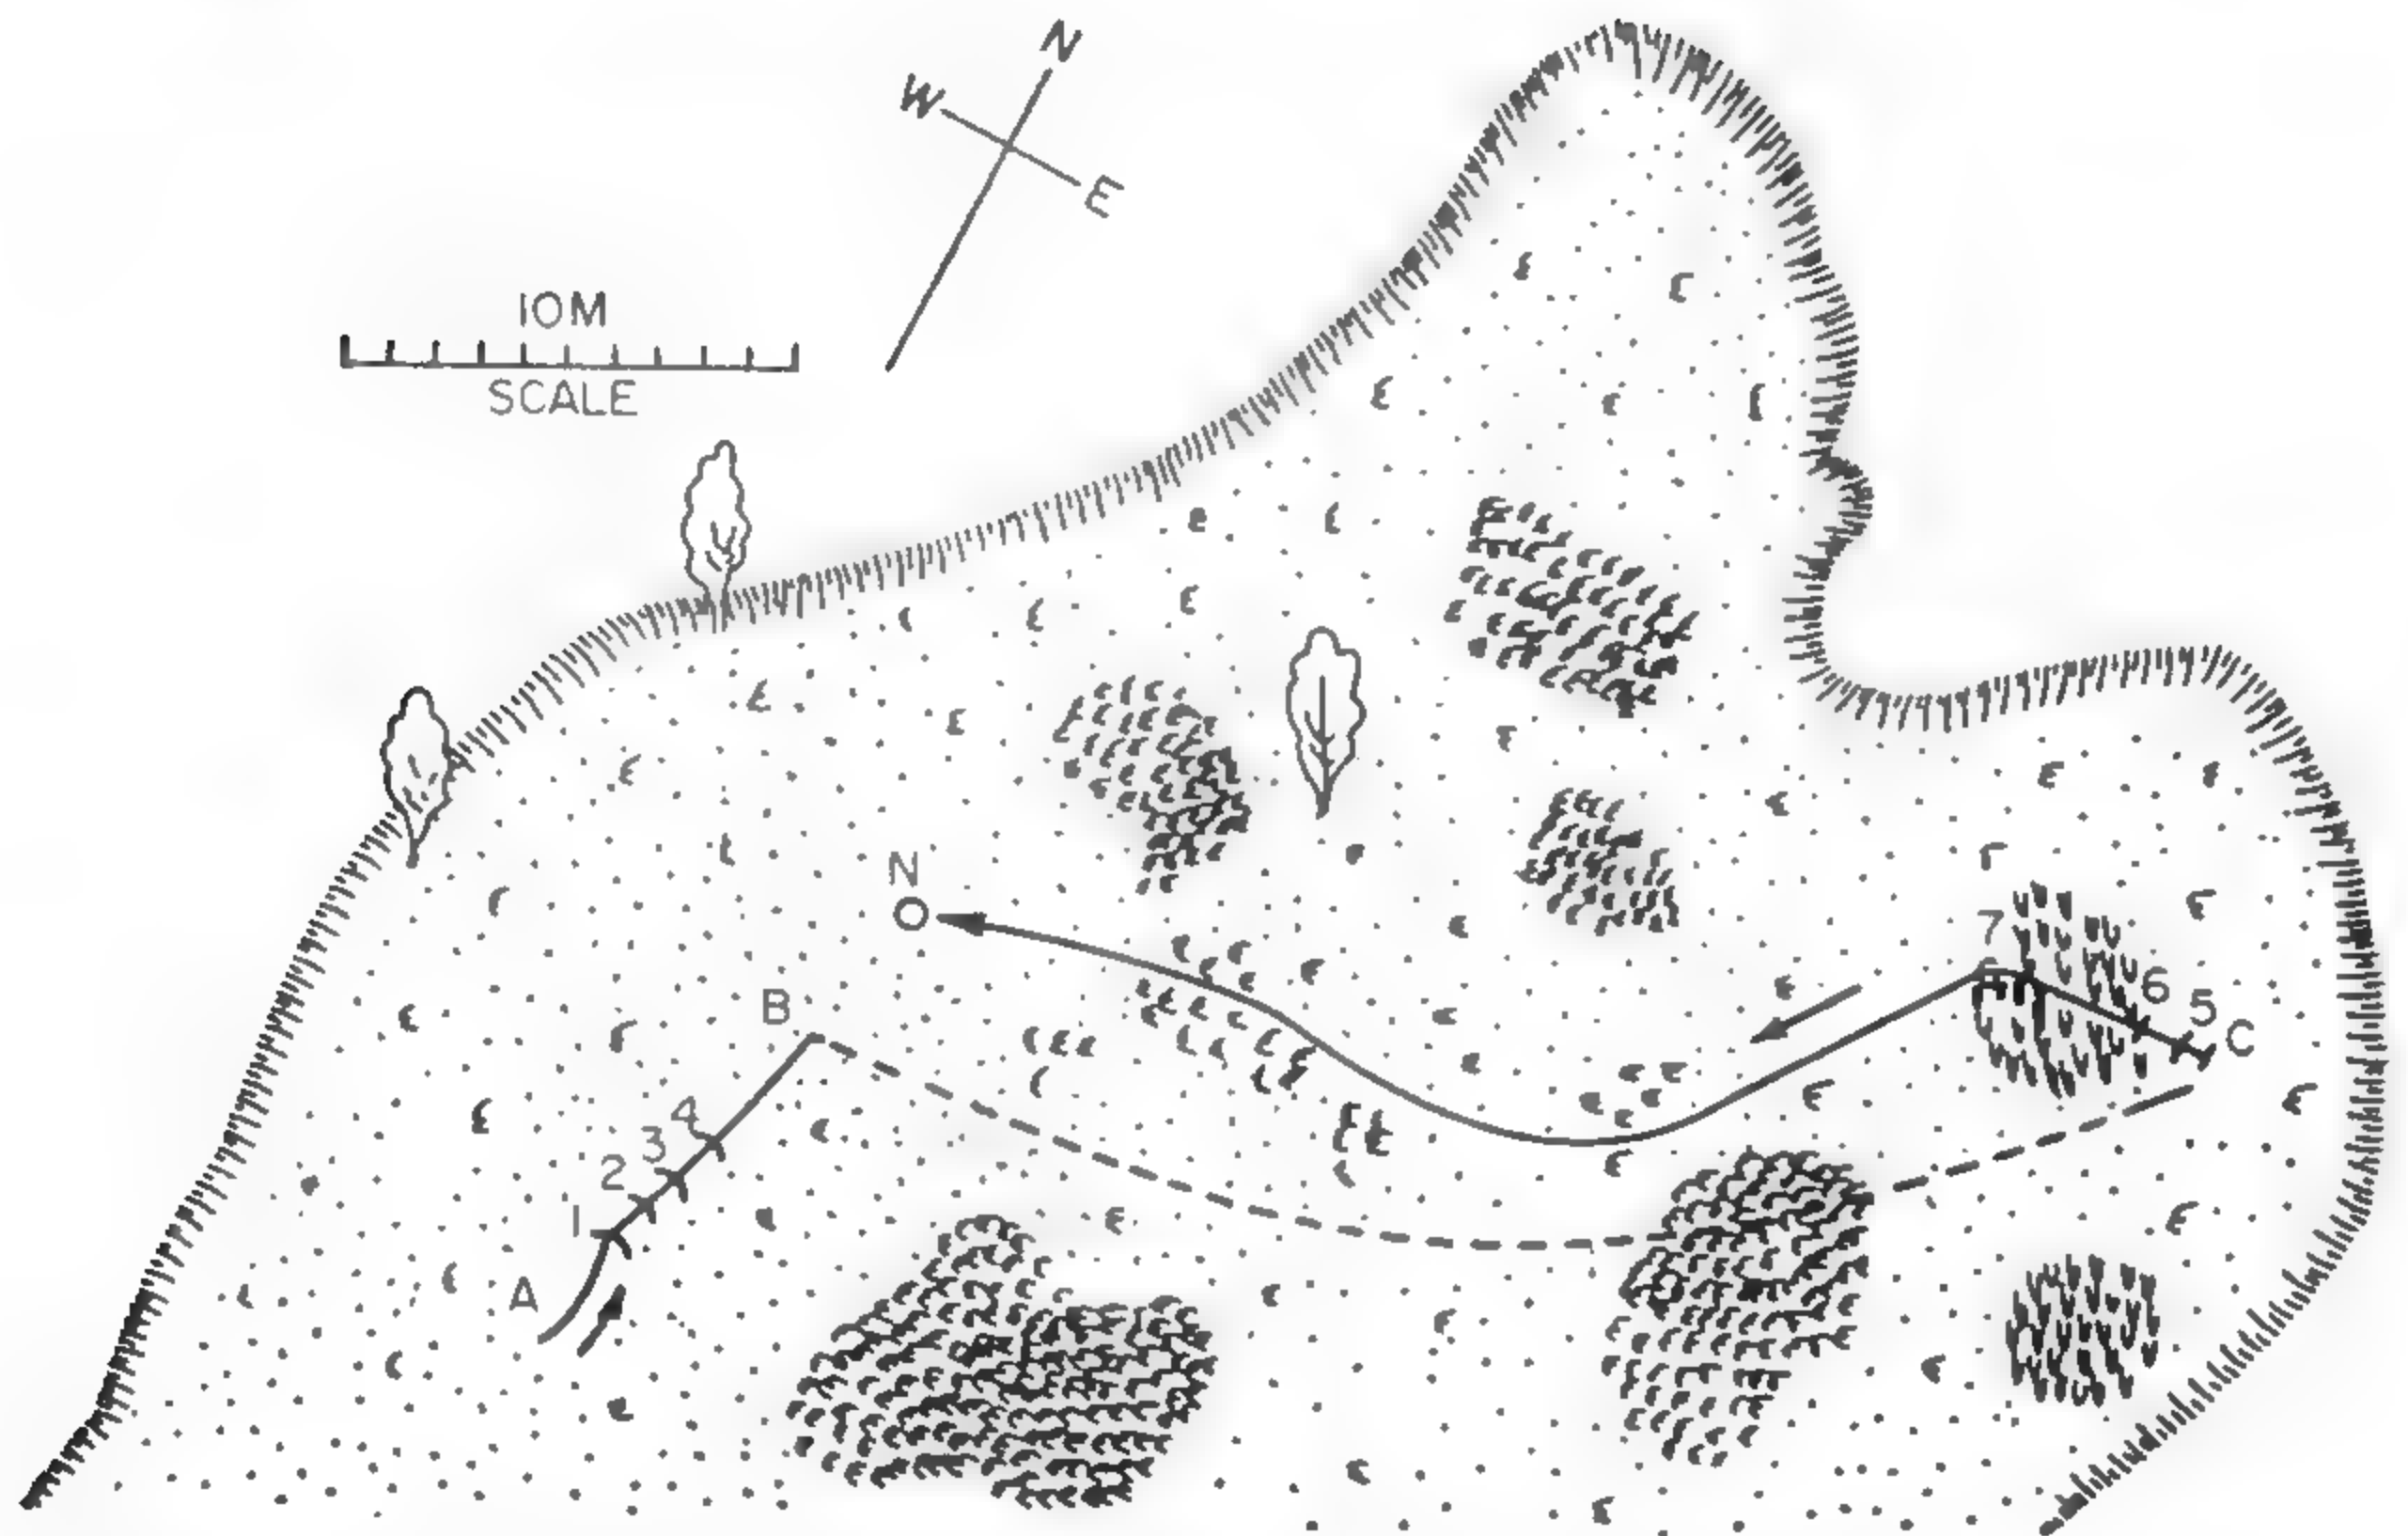


Figure 59. Shows the route taken during detour and displacement experiments on an individual wasp. The heavy line shows the course of the wasp. A is the point at which it was first observed. The numbers 1 to 4 indicate the points at which metal screens were placed in the wasp's path to force detouring. The insect was captured at B. It was transported over the dotted line and released at point C. The remaining numbers indicate further detour experiments. N is the nest, finally reached by the insect. The time taken by the insect was approximately 15 minutes. (Thorpe, 1956.)

weaving of instinctive and learned components. Once an instinctive act has been released, it has opened the way for possible modification. Certain kinds of learning depend upon instinctlike mechanisms (imprinting), and so it is safe to say that learning is modified by instinct as much as instinct is modified by learning.

LEARNING APPLIED TO SOCIAL AND PSYCHOLOGICAL PROBLEMS

We have had little to say—and that only incidentally—about the application of the psychology of learning to general social and psychological problems. More often than not, the worker in the psychology of learning has had in mind only remotely, if at all, an eventual application of his work to the more pressing problems of human behavior. Nevertheless, the application has gone forward, sometimes wisely and sometimes not. A few theorists and experimentalists have argued that the justification for the theoretical study of learning lies in its application, and some of these psychologists have enthusiastically applied the experimental evidence and

theories of learning to problems in personality development and social relations.

Unfortunately, there is some confusion about the application of learning theory, because some ambiguity exists in the meaning of learning theory. In the past generation, theories of behavior have come to be identified with theories of learning. Thus, we find B. F. Skinner generally classified as a learning theorist, though he is very little concerned with learning and memory. With considerably more justification he could be called an incentive theorist. We shall try to avoid this identification of learning with the general theoretical principles of behavior and try to restrict our discussion to associative problems. Of course, we shall not be able to ignore entirely motivational and emotional problems, since these frequently turn out to be based on associations between behavior and its consequences.

The applications of the psychology of learning are numerous, and we can discuss only a few representative examples from the fields of personality, social development, and educational psychology. These should give us some indication of the direction in which applications of learning have gone.

Problems in Personality and Social Development

In recent years there have been many attempts to apply the basic data and theories of learning to the problems of human personality. One of the earliest is found in Guthrie's *Psychology of Human Conflict* (1938). Guthrie attempted to account for problems in personality and personal conflict by relying almost exclusively on classical conditioning as the model for all kinds of behavior, reflex and instrumental, emotional and intellectual. Guthrie's attempt is distinguished by ingenuity and style, but it is more of a tour de force than a serious attempt to cope with the complicated problems of human behavior in terms of associative learning.

More recently, learning theorists have attempted to deal with more specific problems, such as the course of psychotherapy. Furthermore, these attempts, unlike Guthrie's, make some verifiable predictions about the problems with which they deal. In the next few pages we shall examine some of these applications of the basic learning principles to the analysis of personality and psychotherapy.

One of the most persistent efforts of the immediate past has been to achieve a synthesis of psychoanalytic and learning theories. This has intriguing possibilities, and since, as has often been pointed out, Freud was basically an associationist, it has not been so difficult a task as might be

supposed. Much of the work we shall discuss has been written from this synthetic psychoanalytic-learning point of view.

Personality and Psychotherapy

Shoben (1949) suggested that one of the major goals of psychotherapy is the acquisition of symbolic control over behavior. The process of psychotherapy is described by Shoben as the lifting of repression and the development of insight through symbolic reinstatement of the stimuli for anxiety, the diminution of anxiety by counter conditioning, and the process of re-education with the help of the therapist. Within this framework are many problems directly related to the principles of learning.

*Dollard and Miller.*³ One of the most convincing attempts to apply the principles of learning to personality and psychotherapy is demonstrated in Dollard and Miller's *Personality and Psychotherapy* (1950). The authors write from the point of view of Hull's general theory of behavior and learning, but save for a few points, they could have written their book from the point of view of the purely empirical study of learning. Much of their book is devoted to an analysis of a case history and the course of psychotherapy in this case. This gives their book an extremely realistic flavor, but it only begins to suggest the possibilities of the analysis of psychotherapy from the point of view of the learning process.

Dollard and Miller devote some space to outlining the basic principles of conditioning and learning. They then show how these principles can account for the acquisition of fears and anxieties in infancy and childhood. They argue, for example, that if a baby becomes very hungry when it is alone and cries and frets while the hunger drive increases, when it is finally fed there will be a large-scale reduction of need. These responses of crying, fretting, etc., are reinforced and thus become the characteristic responses of the child when left alone.

The bulk of their book is devoted to the consideration of psychotherapy as a learning process and to an analysis of the case illustrating the role of mechanisms of learning in psychotherapy. Dollard and Miller start out by characterizing the neurotic person as someone who learns inadequate ways of thinking and dealing with himself. The inability to think about his basic problems may make the neurotic appear to be stupid; the pressures of conflict make him feel miserable. Neurotic activities and psychosomatic symptoms are the ways in which he has learned to alleviate his misery partially.

³ See the section in Chap. 6 on Dollard and Miller.

The business of psychotherapy, according to Dollard and Miller, is to correct the faulty habits the neurotic has learned. The therapist sets up a warm and permissive atmosphere in which the neurotic's anxieties and fears can be symbolically introduced and allowed to extinguish. The therapist arouses memories of anxiety-inducing events in the neurotic's life history, but it is done in the permissive atmosphere of therapy, so that these fears and anxieties are unreinforced and may therefore be extinguished.

One of the problems for the therapist is to attach labels to things which the neurotic has lost the ability to deal with mentally. Words and sentences are attached to emotions formerly repressed. The neurotic can then deal more adequately with his inner feelings and experiences and he is made capable of dealing with his conflicts at an overt, conscious level. This is essentially the process of discrimination, and these two processes, extinction of fear and anxiety and the development of discrimination, are the essential contributions of psychotherapy.

Dollard and his associates (Dollard, Auld, and White, 1953) have done more than simply present the process of psychotherapy as interpreted by learning theory; they have suggested techniques for the evaluation of therapy and the analysis of its progress based upon this interpretation. Thus, such an approach to psychotherapy is of considerable potential value to the practical clinical psychologist.

Mowrer. A similar analysis of personality disorder and psychotherapy has been presented by Mowrer.⁴ More recently, however, Mowrer has expanded his point of view to encompass a vast multitude of scientific and moral problems in the development of personality and psychotherapy. He has considerably broadened the concepts of learning to be applied to the analysis and description of personality. He makes much more use of a concept of mediational activity (or "pure stimulus act," as Hull has called it) than do Dollard and Miller. For example, he sees thinking as a form of activity in which the organism makes symbolic responses to find out what effects will occur if the action is really carried out (Mowrer, 1954). Thought, says Mowrer, is preparation for action, fantasy is substitution for action. Therefore, although he finds a role for the habit processes much like those suggested by Shoben and by Dollard and Miller, he finds the behavior of the individual in psychotherapy much more complicated than learning theory allows.

Indeed, learning theory does offer a most mechanical and limited model of the individual in difficulty. Since some model is better than none at all,

⁴ See, for example, Mowrer (1950)

Mowrer and others of his general viewpoint run the danger of losing the precision and clarity of a simple model by trying, at this stage, to schematize the most subtle aspects of personality and its development. It is perhaps good that learning models of personality action and psychotherapy are rather simple, for they encourage us to concentrate on solving a few elementary problems. The richer, more luxuriant theories of personality, on the other hand, so overwhelm us with the complexities of the problems to be dealt with, that they discourage investigation. Thus, one of the real contributions of learning to the study of personality should be the simplification it introduces.

Learning and Socialization

Children must learn to become adequate adults in the different societies to which they belong. Societies differ from one another considerably in the extent and kinds of demands they make upon individuals. Consequently, cross-cultural studies of socialization in different societies present an important challenge to the psychology of learning.

Unfortunately, since such a study is just beginning, we cannot expect to find well worked out theoretical applications of the principles of learning. The work of Whiting and Child (1953) has suggested ways of studying the influence of infant and child care on the characteristics of adults in different societies.

Child (1954) points out that the learning of infants and children is more broadly generalized than that of older individuals. This is because young children are preverbal and thus do not have the mediating responses for performing fine discriminations. Such behavior is less susceptible to extinction and verbal control. Secondly, Child points out, the infant is helpless and hence is often subjected to drives of great intensity. Therefore, drive reduction is a basic and powerful element in the control of behavior of very young infants. Thus the stage is set for the learning of powerful emotional reactions and the setting up of secondary drives. The problems of early learning, such as those encountered in toilet training, weaning, and control of aggression, should be powerful influences in shaping the personality of individuals.

Seward (1954) attempts a rigorous derivation from basic learning principles of the role of reinforcement in developing identification in children. Identification is the process of learning the values of others, particularly those of the parent. Seward deduces that the kind of identification developed by punishment will be closely linked to the kinds of behavior actually

punished, while approval will tend to produce identification with much wider generalization.

The role of reinforcement is only one of the many mechanisms that work in the kind of socialization that leads to identification, but it is an important one. Notions like Seward's have implications for a cross-cultural analysis of socialization, for it is possible to examine cultures which differ in the way in which identification is induced to see what differences are produced in the kinds of identification achieved.

Learning and Theories of Personality

It is encouraging to see that the general concepts which have arisen from the study of learning are now being incorporated into the body of personality theory. This is partly because of the close parallel between certain classical psychoanalytic notions and some basic concepts in learning. Over and above this, however, there has been a genuinely new appreciation of the value of some of the familiar ideas in the study of learning for the understanding of personality development.

Rotter (1954), for example, has presented a rather detailed theory of personality development from the point of view of social learning. He adds a new dimension to the understanding of personality by means of a concept of reinforcement value. He reminds us that not only are different reinforcements effective with different people, but that for any individual or group of individuals there is an organized hierarchy of values of reinforcement. Furthermore, these hierarchies evolve with experience, and this kind of change is one aspect of personality development.

Rotter shows that the reinforcement value of a certain reward is to a considerable extent determined by the individual's anticipation of the pattern or frequency of future reward. This notion allows Rotter to reinterpret the role of delay of reinforcement in maladjustive behavior. Various psychologists have supposed that maladjusted individuals, criminals, and neurotics are marked by an inability to act on the basis of delayed reinforcements; such people seek immediate gratification. Thus, says Rotter, if we were to offer a group of children the choice of a penny candy today or a 5 cent candy tomorrow, at least some of the children would choose the penny candy today. From Rotter's point of view, such a choice is not because the delayed reinforcement has a lesser value, but because some individuals have learned that reinforcements promised for the future are not likely to occur.

Rotter and other theorists have proposed a large number of mechanisms

derived from basic studies of learning to account for various aspects of personality development. These concepts should have the effect of leading to fruitful experimental investigation of personality mechanisms.

Problems in Education

At the beginning of the twentieth century there was much hope that the experimental study of learning would lead, in a fairly short time, to a scientific analysis of the problems of school and teaching. Many early books on learning were greatly concerned with the applications of existing facts to problems in education. In a very real sense, some valuable contributions were made to the practice of education. The early studies of transfer, particularly those of Thorndike and Woodworth (1901), were influential in the decline of the notion of formal discipline. All in all, however, the experimental study of learning did not live up to the early hopes for application that it aroused. By 1930, workers in this field disclaimed any attempt to achieve results which would have educational significance. At the same time a great interest in theories of learning arose, and these were mostly concerned with the behavior of animals.

Early Problems in Educational Psychology

The influence of E. L. Thorndike sat heavily upon educational psychology for many years. Thorndike's particular interpretation of learning—which he later called connectionism—provided the core for the account of learning to be found in early textbooks on educational psychology. The principle of effect and trial-and-error behavior were the two central themes in early theories of learning.

There is nothing wrong with an account of reinforcement and trial-and-error processes in a textbook in educational psychology, except that these things by themselves do not make a complete and convincing picture of the learning of children. One of the outcomes of the disappointment with the classical Thorndikian point of view was that a formal discussion of the psychology of learning all but disappeared from accounts of educational psychology. Theories of learning contained so little of interest to educators that such theories were relegated to an extremely minor role in the psychology of education. Since theories of learning had little or nothing of practical value to contribute, their place in the textbooks was taken by new material on personality development, mental hygiene, etc.

Of course, the basic applied material on learning remained. This was mainly an account of the conditions which promote efficient classroom

learning. The textbooks discussed massed and distributed practice, the effect of recitation on memory, and motives in classroom learning. There was little of theoretical interest discussed; the effect of recitation on memory was simply presented as a fact, and its interpretation in terms of recoding and transfer was ignored.

Transfer of training (although not always Thorndike's concept of identical elements) remained as the most basic contribution of learning to educational psychology. This emphasis gradually changed, however. The earlier use of transfer was to further the social-utility movement in education. The apparent lack of transfer from the more traditional school subjects to the real and important concerns of human life was sufficient reason for decreasing the emphasis upon these subjects. The curriculum was revised so that the traditional subjects would be more likely to transfer to the problems of daily living.

More recently, we have come to realize that maximizing transfer is sometimes less important than other considerations. It is perhaps more valuable to produce a limited amount of transfer from some kind of activity that can go on *only* in the schools than to produce a lot of transfer that duplicates the activity of other agencies of the community. Consequently, the emphasis has changed from one of revising the curriculum to maximize transfer to one of maximizing transfer *within* the framework of a curriculum determined by the needs of the schools.

Current Problems in Educational Psychology

Much of educational psychology today is concerned with problems of personality development and individuality. Since there has been a growing interest in the application of theories of learning to individual development and socialization, much of the fundamental theoretical and experimental work in the psychology of learning is now returning to educational psychology, albeit with a different emphasis.

In addition, of course, there have been many developments within educational psychology itself. The dissatisfaction with the traditional material in learning and the conflict between the experimentalists, who insist upon working with pure variables, and the educators, who find that this produces little of direct application, are still with us. A few individuals, however, have been working at bringing together the experimental approach to learning and the problems of classroom learning. Auble and Mech (1953), for example, deal with the application of partial reinforcement to behavior in the classroom. Their results suggest that reinforcement variables must be

treated with extreme caution in applied problems, since children are able to mediate their own reinforcements. The influence of intraverbal processes in children, the nature of the social situation in the classroom, and other conditions make the application of simple variables to classroom learning extremely difficult. The current attention of many experimentalists in classroom learning is now directed toward the study of the interaction between social and personality variables and the variables that control learning.

Despite the many difficulties, Skinner (1954) has recommended a straightforward application of the reinforcement technique for the control of behavior to the classroom. He asks, What reinforcements are available for working with the child in school? The answer is in all of the things—scissors, puzzles, noise makers—which feed back to the child changes in the environment and which are free of aversive, or negatively reinforcing, properties. Skinner argues that in the schoolroom today much of the positive reinforcement from sheer control of the environment is masked by the emotional responses generated by aversive conditions.

Skinner even goes so far as to recommend automatic devices for the control of reinforcement contingencies, since, he argues, the personal control of reinforcement by a teacher is too erratic and irregular. The free use of such devices, says Skinner, would free the teacher to do what she is best suited for—deal with the socialization of the individual child.

Although Skinner's suggestions have outraged many people, some have been impressed. The kinds of techniques he suggests for the control of behavior (and hence of learning) are now undergoing systematic testing. In the next few years we should be able to evaluate them properly.

Recent years have seen little in the way of large-scale theories of educational psychology. Perhaps Skinner's efforts will point in this direction. In a somewhat different direction, Stephens (1951, 1956) has advanced a theory of schooling built on the presupposition that there are many basic nondeliberative (one might almost say instinctive) factors underlying the educative process, and these, more than any particular method or practice in education, are the things that determine how schooling happens.

Stephens started out originally with a hypothesis to account for the large literature of negative findings in the field of experimental comparison of the different "methods" of teaching. These appear to make little or no difference in the outcome of education. This, Stephens argues, is so because different methods can only slightly modify, not introduce or eliminate, some of the fundamental spontaneous factors in education. He argues that such matters as play, discursive conversation, tolerance of unusually communica-

tive people, and the fundamental operations of reinforcement—all work no matter what the specific method. The method supplies direction and emphasis, but the fundamental work is done by the nondeliberative factors. Indeed, these factors are responsible for the institution of schooling itself.

This is radically different from the kind of theory currently predominant in psychological and sociological discussions of education, but it makes much of our knowledge of the pervasiveness of learning as an evolutionary mechanism and of the innately reinforcing properties of nonappetitive rewards, which have received much attention lately. The introduction of such a theory seems to indicate a reawakening of interest in the fundamental connection between learning and educational theories.

IN CONCLUSION

In this short book we have followed a complicated path through the experimental literature on learning. It is disturbing to think that many readers will be so lost in the details that they will fail to see the shape of things as a whole. In these last few paragraphs we shall attempt to give an over-all contour to what we have said.

We have looked upon learning as a primarily associative process. The fundamental variables that control learning are temporal contiguity, similarity (both among responses, implicit and explicit, and among stimuli), repetitions of stimuli and responses, interaction between learned responses, and learned mediating responses. We have seen that the motive-incentive variables such as drive, emotion, and reinforcement (including time and pattern of reinforcement) determine the behavior of organisms. In so far as organisms learn what they do, these variables determine learning. Furthermore, a fundamental class of associations in the animal world is that of responses with reinforcing agencies. We have seen, however, that sometimes when organisms behave without specifically channeled motives or well-defined reinforcements they still learn. Thus, these motive-incentive conditions which guide and control behavior are not essential to learning. The essential features of learning are associations among stimuli and among stimuli and responses as these occur in the experience of the organism.

The variables determining learning interact with species and individual differences among organisms, and they are limited by the capacities of the organisms. We have seen the limitation that the immediate memory span imposes on human memory and how this limitation is circumvented by repetition and recoding. We have seen how intraorganic symbolic processes are

necessary for certain kinds of "insightful" problem solving and how the lack of such processes produces an inevitable limit to the problem-solving ability of certain animals.

We have also discussed the role of language in recoding man's experience. We have had little to say about the learning of linguistic habits themselves. Such habits are, to some extent at least, associative; they are associative in the sense that the learning of nonsense syllables provides a convenient and crude model for the development of such habits. Linguistic habits, however, are more than simply chains of probabilistic associations. Human language has a grammar whose existence makes the learning of natural linguistic habits very different from the learning of nonsense syllables.

Grammar is not purely associationistic. There are certain elementary "forms" in grammar into which we can place words, ideas, even new words and new ideas that have never been used before. We can easily recognize the elements of the common forms of our mother tongue, even though nonsense may be put into those forms. "Gelix frimmaged a leble." This can be identified by most intelligent users of English as a simple, declarative sentence. It is not particularly difficult to identify the subject, the verb, the object, and the syntax in which these appear. Part of our learning of language is learning the sequential and semantic associations that are the common property of the community, but another part is the learning of linguistic forms into which we can place our words.

We have had little or nothing to say about how such grammatical forms are learned and used, for the very good reason that this is one of the problems in verbal learning that has remained almost untouched. It is one of the really few virginal territories left for research in verbal learning, and it is to be recommended to anyone who wishes to make a fundamental contribution to the understanding of the human mind and its verbal processes. Therefore, the great emphasis upon sequential and mediational associative processes we have made in this book needs some correction. These associations *are* important, but associations imposed within grammatical forms are also important, and we have as yet few tools for dealing with them.

In our discussion of verbal learning we had little to say about meaning, mainly because what can be said is not primarily a problem in learning. Suffice it to say that the prevailing view in psycholinguistics today is that meaning has its roots in association. The meaning of a word, a phrase, or any kind of verbal unit is determined by the associations of that unit with other verbal and nonverbal responses and with external stimuli generally. Meanings, indeed, seem frequently to be little more than the bridge between

one associative cluster and another. Thus, the general question of meaning is likely to get lost in the details of recoding, associative clustering, and stimulus generalization.

There are many points at which the study of learning crosses the other disciplines concerned with man's mind and its works. Language, neurophysiology, personality development, and social anthropology—all cross in different ways. If we have achieved some understanding of the way in which organisms learn, we shall carry to these other disciplines something of fundamental importance.

One basic idea that has come out of the systematic study of learning is the attitude that some of the complexities of human life and society can be understood in terms of the operation of basic learning laws. If the implications of this idea have begun to stir, we will have accomplished something.

Always, of course, a discipline such as the psychology of learning promises more than it can achieve. It leaves a model of the organism that is stiff, mechanical, and only a faint copy of reality. It is, however, better than no model at all, and if we are to understand the behavior of organisms in an abstract, scientific sense we badly need such a model. The course of future work in learning and in the basic theory of behavior that goes along with it will be in the direction of introducing some flexibility and life into the model. At the same time, perhaps paradoxically, the model will probably become more abstract. In other words, certain concepts we now use, which are crude physical models, will tend to become more and more abstract and mathematical in character. This dual process of introducing flexibility and abstractness into the behavioral model of organisms appears to be the fundamental task of future research in the psychology of learning.

It also is beginning to appear more likely in fact, rather than merely in hope, that the psychology of learning will begin to approach a unity of concepts with other disciplines. Thus the fundamental associative concepts of the future will probably be stated in such a way that they have some neurophysiological as well as behavioral validity. All this is in the future, however, and at present our horizons are limited by our ability to see only the dimmest connections between learning and the nervous system and between learning and man's best intellectual efforts.

What has gone into this book is a short account of the present status of the psychology of learning. Like all healthy, growing scientific efforts, it is changing continuously. Some of the gaping deficiencies and points of obscurity and downright error will soon be corrected, but it is unlikely that the main outlines will change very much in the near future.

REFERENCES

- ABORN, M. (1953) The influence of experimentally induced failure on the retention of material acquired through set and incidental learning. *J. exp. Psychol.*, **45**, 225-231.
- ABORN, M., AND RUBENSTEIN, H. (1952) Information theory and immediate recall. *J. exp. Psychol.*, **44**, 260-266.
- ADAMS, D. K. (1929) Experimental studies of adaptive behavior in cats. *Comp. Psychol. Monogr.*, **6**, No. 27.
- ADAMSON, R. E. (1952) Functional fixedness as related to problem solving. A repetition of three experiments. *J. exp. Psychol.*, **44**, 288-291.
- ADAMSON, R. E., AND TAYLOR, D. W. (1954) Functional fixedness as related to elapsed time and to set. *J. exp. Psychol.*, **47**, 122-126.
- ADELMAN, H. M., AND MAATSCH, J. L. (1955) Resistance to extinction as a function of the type of response elicited by frustration. *J. exp. Psychol.*, **50**, 61-65.
- ALBERTS, E., AND EHRENFREUND, D. (1951) Transposition in children as a function of age. *J. exp. Psychol.*, **41**, 30-38.
- AMMONS, R. B. (1947) Acquisition of motor skills: I. Quantitative analysis and theoretical formulation. *Psychol. Rev.*, **54**, 263-281.
- ANDERSON, E. (1941a) The externalization of drive: III. Maze learning by non-rewarded and satiated rats. *J. genet. Psychol.*, **59**, 397-426.
- ANDERSON, E. (1941b) The externalization of drive: I. Theoretical considerations. *Psychol. Rev.*, **48**, 204-224.
- ARNOLD, W. J. (1947) Simple reaction chains and their integration: I. Homogeneous chaining with terminal reinforcement. *J. comp. physiol. Psychol.*, **40**, 349-364.
- ARNOULT, M. D. (1956) Familiarity and recognition of nonsense shapes. *J. exp. Psychol.*, **51**, 269-276.
- ATTNEAVE, F. (1954) Some informational aspects of visual perception. *Psychol. Rev.*, **61**, 183-193.
- ATTNEAVE, F. (1955) Symmetry, information and memory for patterns. *Amer. J. Psychol.*, **68**, 209-222.
- ATWATER, S. K. (1953) Proactive inhibition and associative facilitation as affected by degree of prior learning. *J. exp. Psychol.*, **46**, 400-404.
- AUBLE, D., AND MECH, E. V. (1953) Partial verbal reinforcement related to distributed practice in a classroom situation. *J. Psychol.*, **36**, 165-186.
- BAKER, R. A. (1953) Aperiodic feeding behavior in the albino rat. *J. comp. physiol. Psychol.*, **46**, 422-426.
- BAKER, R. A., AND LAWRENCE, D. H. (1951) The differential effects of simultaneous and successive stimuli presentation on transposition. *J. comp. Psychol.*, **44**, 378-382.

- BARTLETT, F. C. (1932) *Remembering: A study in experimental and social psychology*. New York: Cambridge.
- BASS, M. J., AND HULL, C. L. (1934) Irradiation of a tactile conditioned reflex in man. *J. comp. Psychol.*, **17**, 47-65.
- BATTIG, W. F. (1956) Transfer from verbal pretraining to motor performance as a function of motor task complexity. *J. exp. Psychol.*, **51**, 371-378.
- BAUER, F. J., AND LAWRENCE, D. H. (1953) Influence of similarity of choice-point and goal cues on discrimination learning. *J. comp. physiol. Psychol.*, **46**, 241-248.
- BAUM, M. H. (1954) Simple concept learning as a function of intralist generalization. *J. exp. Psychol.*, **47**, 89-94.
- BIRCH, H. G. (1945) The relation of previous experience to insightful problem-solving. *J. comp. Psychol.*, **38**, 367-383.
- BITTERMAN, M. E., AND COATE, W. B. (1950) Some new experiments on the nature of discrimination learning in the rat. *J. comp. physiol. Psychol.*, **43**, 198-210.
- BITTERMAN, M. E., AND HOLTZMAN, W. H. (1952) Conditioning and extinction of the galvanic skin response as a function of anxiety. *J. abnorm. soc. Psychol.*, **47**, 615-623.
- BLODGETT, H. C. (1929) The effect of the introduction of reward upon maze behavior in rats. *Univ. Calif. Publ. Psychol.*, **4**, 113-134.
- BLODGETT, H. C., AND MCCUTCHAN, K. (1947) Place vs. response learning in the simple T-maze. *J. exp. Psychol.*, **37**, 412-422.
- BOOK, W. F. (1925) *Learning to typewrite*. New York: Gregg.
- BOUSFIELD, W. A. (1953) The occurrence of clustering in the recall of randomly arranged associates. *J. gen. Psychol.*, **49**, 229-240.
- BOUSFIELD, W. A., AND COHEN, B. H. (1953) The effects of reinforcement on the occurrence of clustering in the recall of randomly arranged associates. *J. Psychol.*, **36**, 67-81.
- BOUSFIELD, W. A., AND COHEN, B. H. (1955) The occurrence of clustering in the recall of randomly arranged words of different frequencies of usage. *J. gen. Psychol.*, **52**, 83-95.
- BOUSFIELD, W. A., COHEN, B. H., AND SILVA, J. G. (1956) The extension of Marbe's law to the recall of stimulus words. *Amer. J. Psychol.*, **69**, 429-433.
- BRANDAUER, C. M. (1953) A confirmation of Webb's data concerning the action of irrelevant drives. *J. exp. Psychol.*, **45**, 150-152.
- BRIGGS, G. E. (1954) Acquisition, extinction and recovery functions in retroactive inhibition. *J. exp. Psychol.*, **47**, 285-293.
- BRIGGS, G. E. (1957) Retroactive inhibition as a function of the degree of original and interpolated learning. *J. exp. Psychol.*, **53**, 60-67.
- BRIGGS, L. J., AND REED, H. B. (1943) The curve of retention for substance material. *J. exp. Psychol.*, **32**, 513-517.
- BROGDEN, W. J., LIPMAN, E. A., AND CULLER, E. A. (1938) The role of incentive in conditioning and extinction. *Amer. J. Psychol.*, **51**, 109-117.
- BROGDEN, W. J., AND SCHMIDT, R. E. (1954a) Effect of number of choices per unit of a verbal maze on learning and serial position errors. *J. exp. Psychol.*, **47**, 235-240.
- BROGDEN, W. J., AND SCHMIDT, R. E. (1954b) Acquisition of a 24-unit verbal maze as a function of number of alternate choices per unit. *J. exp. Psychol.*, **48**, 335-338.
- BROWN, J. S. (1948) Gradients of approach and avoidance responses and their relation to motivation. *J. comp. physiol. Psychol.*, **41**, 450-465.
- BROWN, J. S. (1953) Problems presented by the concept of acquired drives. In Jones, M. R. (Ed.) *Current theory and research in motivation, a Symposium*. Lincoln, Nebr.: Univer. of Nebraska Press.

- BROWN, J. S., BILODEAU, E. A., AND BARON, M. R. (1951) Bidirectional gradients in the strength of a generalized voluntary response to stimuli on a visual-spatial dimension. *J. exp. Psychol.*, **41**, 52-61.
- BRUCE, R. W. (1933) Conditions of transfer of training. *J. exp. Psychol.*, **16**, 343-361.
- BRUNER, J. S., GOODNOW, J. J., AND AUSTIN, G. A. (1956) *A study of thinking*. New York: Wiley.
- BRUNSWIK, E. (1939) Probability as a determiner of rat behavior. *J. exp. Psychol.*, **25**, 175-197.
- BRUNSWIK, E. (1956) *Perception and the representative design of psychological Experiments*. Berkeley, Calif.: Univer. of Calif. Press.
- BRYAN, W. L., AND HARTER, N. (1897) Studies in the physiology and psychology of the telegraphic language. *Psychol. Rev.*, **4**, 27-53.
- BRYAN, W. L., AND HARTER, N. (1899) Studies on the telegraphic language: the acquisition of a hierarchy of habits. *Psychol. Rev.*, **6**, 345-375.
- BUCHWALD, A. M., AND YAMAGUCHI, H. G. (1955) The effect of change in drive level on habit reversal. *J. exp. Psychol.*, **50**, 265-268.
- BUGELSKI, B. R. (1942) Interference with recall of original responses after learning new responses to old stimuli. *J. exp. Psychol.*, **30**, 368-379.
- BUGELSKI, B. R., AND CADWALLADER, T. C. (1956) A reappraisal of the transfer and retroaction surface. *J. exp. Psychol.*, **52**, 360-366.
- BUGELSKI, B. R., COYER, R. A., AND ROGERS, W. A. (1952) A criticism of pure acquisition and pre-extinction expectancies. *J. exp. Psychol.*, **44**, 27-30.
- BUNCH, M. E., AND MAGSDICK, W. K. (1933) The retention in rats of an incompletely learned maze solution for short intervals of time. *J. comp. Psychol.*, **16**, 385-409.
- BURTT, H. E. (1941) An experimental study of early childhood memory. *J. genet. Psychol.*, **58**, 435-439.
- BUSH, R. R., AND MOSTELLER, F. (1951) A model for stimulus generalization and discrimination. *Psychol. Rev.*, **58**, 413-423.
- BUSH, R. R., AND MOSTELLER, F. (1955) *Stochastic models for learning*. New York: Wiley.
- BUSS, A. H., WEINER, M., AND BUSS, E. (1954) Stimulus generalization as a function of verbal reinforcement combinations. *J. exp. Psychol.*, **48**, 433-436.
- BUTLER, R. A. (1954) Incentive conditions which influence visual exploration. *J. exp. Psychol.*, **48**, 19-23.
- BUTLER, R. A., AND HARLOW, H. F. (1954) Persistence of visual exploration in monkeys. *J. comp. physiol. Psychol.*, **47**, 258-263.
- BUXTON, C. E. (1940) Latent learning and the goal gradient hypothesis. *Contr. Psychol. Theor.*, **2**, No. 6.
- CALVIN, A. D., CLIFFORD, L. T., CLIFFORD, B., BOLDEN, L., AND HARVEY, J. (1956) Experimental validation of conditioned inhibition. *Psychol. Rep.*, **2**, 51-56.
- CAMPBELL, B. A., AND KRAELING, D. (1954) Response strength as a function of drive level during training and extinction. *J. comp. physiol. Psychol.*, **47**, 101-103.
- CAMPBELL, B. A., AND SHEFFIELD, F. D. (1953) Relation of random activity to food deprivation. *J. comp. physiol. Psychol.*, **46**, 320-322.
- CARPER, J. W. (1953) A comparison of the reinforcing value of a nutritive and non-nutritive substance under conditions of specific and general hunger. *Amer. J. Psychol.*, **66**, 270-277.
- CARR, H. A., AND WATSON, J. B. (1908) Orientation in the white rat. *J. comp. Neurol.*, **18**, 27-44.
- CHILD, I. L. (1954) Socialization. In Lindzey, G. (Ed.) *Handbook of social psychology*, Vol. 2. Cambridge, Mass.: Addison-Wesley.

- COFER, C. N. (1941) A comparison of logical and verbatim learning of prose passages of different lengths. *Amer. J. Psychol.*, **54**, 1-20.
- COFER, C. N. (1954) The role of language in human problem solving. Paper presented to Conference on Human Problem Solving, New York Univer.
- COFER, C. N., AND FOLEY, J. P., JR. (1942) Mediated generalization and the interpretation of verbal behavior. I. Prolegomena. *Psychol. Rev.*, **49**, 513-540.
- COLE, L. E. (1924) Latin as a preparation for French and Spanish. *Sch. & Soc.*, **19**, 618-622.
- COTTERMAN, T. E., MEYER, D. R., AND WICKENS, D. D. (1956) Discrimination reversal learning in Marmosets. *J. comp. physiol. Psychol.*, **49**, 539-541.
- COWLES, J. T. (1937) Food tokens as incentives for learning by chimpanzees. *Comp. Psychol. Monogr.*, **14**, No. 7.
- CRAWFORD, M. P. (1947) Psychological research on operational training in the continental air forces. *AAF Aviation Psychol. Program Res. Rep.*, No. 16.
- CRESPI, L. P. (1942) Quantitative variation of incentive and performance in the white rat. *Amer. J. Psychol.*, **55**, 467-517.
- CRESPI, L. P. (1944) Amount of reinforcement and level of performance. *Psychol. Rev.*, **51**, 341-357.
- DATTMAN, P. E., AND ISRAEL, H. (1951) The order of dominance among conceptual capacities: an experimental test of Heidebreder's hypothesis. *J. Psychol.*, **31**, 147-160.
- DEESE, J. (1950) A quantitative derivation of latent learning. *Psychol. Rev.*, **57**, 291-294.
- DEESE, J. (1951) Extinction of a discrimination without performance of the choice response. *J. comp. physiol. Psychol.*, **44**, 362-366.
- DEESE, J. (1956) Complexity of contour in the recognition of visual form. *WADC Tech. Rep.*, No. 56-60.
- DEESE, J., AND HARDMAN, G. W., JR. (1954) An analysis of errors in retroactive inhibition of rote verbal learning. *Amer. J. Psychol.*, **67**, 299-307.
- DEESE, J., AND CARPENTER, J. A. (1951) Drive-level and reinforcement. *J. exp. Psychol.*, **42**, 236-238.
- DEESE, J., AND KAUFMAN, R. A. (1957) Serial effects in recall of unorganized and sequentially organized verbal material. *J. exp. Psychol.*, **54**, 180-187.
- DEESE, J., AND KRESSE, F. H. (1952) An analysis of errors in rote serial learning. *J. exp. Psychol.*, **43**, 199-202.
- DEESE, J., AND LAZARUS, R. S. (1952) The effect of psychological stress upon perceptual-motor performance. *Human Resources Res. Bull.*, No. 52-19.
- DEESE, J., LAZARUS, R. S., AND KEENAN, J. (1953) Anxiety, anxiety reduction and stress in learning. *J. exp. Psychol.*, **46**, 55-60.
- DEESE, J., AND MARDER, V. J. (1957) The pattern of errors in delayed recall of serial learning after interpolation. *Amer. J. Psychol.*, **70**, 594-599.
- DETHIER, V. G. (1957) Communication by insects. *Science*, **125**, 331-335.
- DINSMOOR, J. A. (1952) The effect of hunger on discriminated responding. *J. abnorm. soc. Psychol.*, **47**, 67-72.
- DINSMOOR, J. A. (1954) Punishment: I. The avoidance hypothesis. *Psychol. Rev.*, **61**, 34-46.
- DOLLARD, J., AULD, F., AND WHITE, A. M. (1953) *Steps in psychotherapy*. New York: Macmillan.
- DOLLARD, J., AND MILLER, N. E. (1950) *Personality and psychotherapy*. New York: McGraw-Hill.

- DUFORT, R. H., AND KIMBLE, G. A. (1956) Changes in response strength with changes in the amount of reinforcement. *J. exp. Psychol.*, **51**, 185-191.
- DUNCAN, C. P. (1951) The effect of unequal amounts of practice on motor learning before and after rest. *J. exp. Psychol.*, **42**, 257-264.
- DUNCAN, C. P., AND UNDERWOOD, B. J. (1953) Retention of transfer in motor learning after twenty-four hours and after fourteen months. *J. exp. Psychol.*, **46**, 445-452.
- DUNCKER, K. (1945) On problem solving. (Trans. from 1935 original.) *Psychol. Monogr.*, **58**, No. 270.
- DYKMAN, R. A., GANTT, W. H., AND WHITEHORN, J. C. (1956) Conditioning as emotional sensitization and differentiation. *Psychol. Monogr.*, **70**, No. 15.
- EBBINGHAUS, H. (1885) *Über das Gedächtnis: Untersuchungen zur experimentellen Psychologie*. Leipzig: Duncker und Humblot.
- ECCLES, J. C. (1953) *The neurophysiological basis of mind*. New York: Oxford.
- ECCLES, J. C., AND RALL, W. (1951) Effects induced in a monosynaptic path by its activation. *J. neurophysiol.*, **14**, 353-376.
- EHRENFREUND, D. (1948) An experimental test of the continuity theory of discrimination learning with pattern vision. *J. comp. physiol. Psychol.*, **41**, 408-422.
- EHRENFREUND, D. (1952) A study of the transposition gradient. *J. exp. Psychol.*, **43**, 83-87.
- ELLIS, D. S. (1953) Inhibition theory and the effort variable. *Psychol. Rev.*, **60**, 383-392.
- ELLSON, D. G. (1938) Quantitative studies of the interaction of simple habits: I. Recovery from specific and generalized effects of extinction. *J. exp. Psychol.*, **23**, 339-358.
- ERIKSEN, C. W. (1954) Some personality correlates of stimulus generalization under stress. *J. abnorm. soc. Psychol.*, **49**, 561-565.
- ESTES, K. W. (1944) An experimental study of punishment. *Psychol. Monogr.*, **57**, No. 263.
- ESTES, K. W. (1950) Towards a statistical theory of learning. *Psychol. Rev.*, **57**, 94-107.
- ESTES, K. W. (1955) Statistical theory of spontaneous recovery. *Psychol. Rev.*, **62**, 145-154.
- ESTES, K. W. (1956) The problems of inference from curves based on group data. *Psychol. Bull.*, **53**, 134-140.
- ESTES, K. W., AND BURKE, C. J. (1953) A theory of stimulus variability in learning. *Psychol. Rev.*, **60**, 276-286.
- ESTES, K. W., AND STRAUGHAN, J. H. (1954) Analysis of a verbal conditioning situation in terms of statistical learning theory. *J. exp. Psychol.*, **47**, 225-234.
- FARBER, I. E., AND SPENCE, K. W. (1953) Complex learning and conditioning as a function of anxiety. *J. exp. Psychol.*, **45**, 120-125.
- FELDMAN, R. S. (1953) The specificity of the fixated response in the rat. *J. comp. physiol. Psychol.*, **46**, 487-492.
- FINAN, J. L. (1940) Quantitative studies in motivation. I. Strength of conditioning in rats under various degrees of hunger. *J. comp. Psychol.*, **29**, 119-124.
- FINGER, F. W. (1949) Hunger drive and general activity in the rat. *Amer. Psychologist*, **4**, 223-224.
- FRANDSEN, A. (1957) *How children learn*. New York: McGraw-Hill.
- FRENCH, E. G. (1955) Some characteristics of achievement motivation. *J. exp. Psychol.*, **50**, 232-236.
- FRENCH, R. S. (1953) Number of common elements and consistency of reinforcement in a discrimination learning task. *J. exp. Psychol.*, **45**, 25-34.

- FRENCH, R. S. (1954) Pattern recognition in the presence of visual noise. *J. exp. Psychol.*, **47**, 22-26.
- FREUD, S. (1925) *Collected papers*. Vol. 4. London: Hogarth.
- FRICK, F. C. (1953) The effect of anxiety—a problem in measurement. *J. comp. physiol. Psychol.*, **46**, 120-123.
- FRICK, F. C., AND MILLER, G. A. (1951) A statistical description of operant conditioning. *Amer. J. Psychol.*, **64**, 20-36.
- FURCHTGOTT, E., AND RUBIN, R. D. (1953) The effect of magnitude of reward on maze learning in the white rat. *J. comp. physiol. Psychol.*, **46**, 9-12.
- GAGNE, R. M. (1941) The effect of spacing of trials on the acquisition and extinction of a conditioned operant response. *J. exp. Psychol.*, **29**, 201-216.
- GAGNE, R. M., AND BAKER, K. E. (1950) Stimulus pre-differentiation as a factor in transfer of training. *J. exp. Psychol.*, **40**, 439-451.
- GAGNE, R. M. (1950) Stimulus pre-differentiation as a factor in transfer of training. *J. exp. Psychol.*, **40**, 439-451.
- GAGNE, R. M., FOSTER, H., AND CROWLEY, M. E. (1948) The measurement of transfer of training. *Psychol. Bull.*, **45**, 97-130.
- GALANTER, E., AND SHAW, W. A. (1954) "Cue" vs. "reactive inhibition" in place and response learning. *J. comp. physiol. Psychol.*, **47**, 395-398.
- GANTT, W. H. (1936) An experimental approach to psychiatry. *Amer. J. Psychiat.*, **92**, 1007-1021.
- GATLING, F. (1952) The effect of repeated stimulus reversals on learning in the rat. *J. comp. physiol. Psychol.*, **45**, 347-351.
- GIBSON, E. J. (1941) Retroactive inhibition as a function of degree of generalization between tasks. *J. exp. Psychol.*, **28**, 93-115.
- GLADSTONE, A. E. (1948) Reactively homogeneous compound trial and error learning with distributed trials and serial reinforcement. *J. exp. Psychol.*, **38**, 289-297.
- GLANZER, M. (1953) The role of stimulus satiation in spontaneous alternation. *J. exp. Psychol.*, **45**, 387-393.
- GLAZE, J. A. (1928) The association value of nonsense syllables. *J. genet. Psychol.*, **35**, 255-267.
- GLEITMAN, H., NACHMIAS, J., AND NEISSER, U. (1954) The S-R reinforcement theory of extinction. *Psychol. Rev.*, **61**, 23-33.
- GODBEER, E. (1940) Factors inducing conflict in the choice behavior of children. New Haven, Conn., Dissertation, Yale Univer. Cited in Williams (1943).
- GOODWIN, W. R., AND LAWRENCE, D. H. (1955) The functional independence of two discrimination habits associated with a constant stimulus situation. *J. comp. physiol. Psychol.*, **48**, 437-443.
- Goss, A. E. (1953) Transfer as a function of type and amount of preliminary experience with task stimuli. *J. exp. Psychol.*, **46**, 419-428.
- GRAHAM, C. H., AND GAGNE, R. M. (1940) The acquisition, extinction and spontaneous recovery of a conditioned operant response. *J. exp. Psychol.*, **26**, 251-280.
- GRANT, D. A., HAKE, H. W., AND HORNSETH, J. P. (1951) Acquisition and extinction of a verbal conditioned response with differing percentages of reinforcement. *J. exp. Psychol.*, **42**, 1-5.
- GRANT, D. A., AND SCHILLER, J. J. (1953) Generalization of the conditioned galvanic skin response to visual stimuli. *J. exp. Psychol.*, **46**, 309-313.
- GRANT, D. A., SCHIPPER, L. M., AND ROSS, B. M. (1952) Effect of intertrial interval during acquisition and extinction of the conditioned eyelid response following partial reinforcement. *J. exp. Psychol.*, **44**, 203-210.
- GREENSPOON, J. (1955) The reinforcing effect of two spoken sounds on the frequency of two responses. *Amer. J. Psychol.*, **68**, 409-416.

- GRICE, G. R. (1948) The relation of secondary reinforcement to delayed reward in visual discrimination learning. *J. exp. Psychol.*, **38**, 1-16.
- GRICE, G. R. (1949) Visual discrimination learning with simultaneous and successive presentation of stimuli. *J. comp. physiol. Psychol.*, **42**, 365-373.
- GRICE, G. R. (1951) The acquisition of a visual discrimination habit following extinction of response to one stimulus. *J. comp. physiol. Psychol.*, **44**, 149-153.
- GRICE, G. R., AND SALTZ, E. (1950) The generalization of an instrumental response to stimuli varying in the size dimension. *J. exp. Psychol.*, **40**, 702-708.
- GRINDLEY, G. C. (1929-1930) Experiments on the influence of amount of reward on learning in young chickens. *Brit. J. Psychol.*, **20**, 173-180.
- GULLIKSEN, H. (1934) A rational equation of the learning curve based on Thorndike's law of effect. *J. gen. Psychol.*, **11**, 395-434.
- GUTHRIE, E. R. (1935) *The psychology of learning*. New York: Harper.
- GUTHRIE, E. R. (1938) *The psychology of human conflict*. New York: Harper.
- GUTHRIE, E. R. (1952) *The psychology of learning*. (Rev. ed.) New York: Harper.
- GUTHRIE, E. R., AND HORTON, G. P. (1946) *Cats in a puzzle box*. New York: Rinehart.
- GUTTMAN, N. (1953) Operant conditioning, extinction and periodic reinforcement in relation to concentration of sucrose used as a reinforcing agent. *J. exp. Psychol.*, **46**, 213-224.
- HARLOW, H. F. (1949) The formation of learning sets. *Psychol. Rev.*, **56**, 51-65.
- HARLOW, H. F. (1950a) Learning and satiation of response in intrinsically motivated complex puzzle performance by monkeys. *J. comp. physiol. Psychol.*, **43**, 289-294.
- HARLOW, H. F. (1950b) Analysis of discrimination learning by monkeys. *J. exp. Psychol.*, **40**, 26-39.
- HARLOW, H. F. (1951) Thinking. In Helson, H. (Ed.) *Theoretical foundations of Psychology*. Princeton, N. J.: Van Nostrand.
- HARLOW, H. F., AND MCCLEARN, G. E. (1954) Object discriminations learned by monkeys on the basis of manipulation motives. *J. comp. physiol. Psychol.*, **47**, 71-74.
- HARLOW, H. F., AND MEYER, D. (1952) Paired-comparison scales for monkey rewards. *J. comp. physiol. Psychol.*, **45**, 73-79.
- HEATHERS, G. L. (1940) The avoidance of repetition of a maze reaction in the rat as a function of the time interval between trials. *J. Psychol.*, **10**, 359-380.
- HEBB, D. O. (1949) *The organization of behavior*. New York: Wiley.
- HEIDBREDER, E. (1924) An experimental study of thinking. *Arch. Psychol.*, N.Y., **11**, No. 73.
- HEIDBREDER, E. (1946a) The attainment of concepts: I. Terminology and methodology. *J. gen. Psychol.*, **35**, 173-189.
- HEIDBREDER, E. (1946b) The attainment of concepts: II. The problem. *J. gen. Psychol.*, **35**, 191-223.
- HEIDBREDER, E. (1947) The attainment of concepts: III. The process. *J. gen. Psychol.*, **24**, 93-138.
- HILGARD, E. R. (1956) *Theories of learning*. (2d ed.) New York: Appleton-Century-Crofts.
- HILGARD, E. R., AND MARQUIS, D. M. (1940) *Conditioning and learning*. New York: Appleton-Century-Crofts.
- HILLMAN, B., HUNTER, W. S., AND KIMBLE, G. A. (1953) The effect of drive level on the maze performance of the white rat. *J. comp. physiol. Psychol.*, **46**, 87-89.
- HINDE, R. A. (1954) Changes in responsiveness to a constant stimulus. *Brit. J. anim. Behav.*, **2**, 41-55.
- HOBBS, N. (1947) Psychological research on flexible gunnery training. *AAF Aviation Psychology Program Res. Rep. No. 11*.

- HOLLAND, G. A. (1953) Simple trial-and-error learning: massed and distributed trials with habits of unequal initial strength. *J. comp. physiol. Psychol.*, **46**, 90-94.
- HOMME, L. E. (1956) Spontaneous recovery and statistical learning. *J. exp. Psychol.*, **51**, 205-212.
- HOVLAND, C. I. (1936) "Inhibition of reinforcement" and phenomena of experimental extinction. *Proc. Nat. Acad. Sci. Wash.*, **22**, 430-433.
- HOVLAND, C. I. (1937a) The generalization of conditioned responses: I. The sensory generalization of conditioned responses with varying frequencies of tone. *J. gen. Psychol.*, **17**, 125-148.
- HOVLAND, C. I. (1937b) The generalization of conditioned responses: II. The sensory generalization of conditioned responses with varying intensities of tone. *J. genet. Psychol.*, **51**, 279-291.
- HOVLAND, C. I. (1938) Experimental studies in rote-learning theory: II. Reminiscence with varying speeds of syllable presentation. *J. exp. Psychol.*, **22**, 338-353.
- HOVLAND, C. I. (1940) Experimental studies in rote-learning theory: VII. Distribution of practice with varying lengths of list. *J. exp. Psychol.*, **27**, 271-284.
- HOVLAND, C. I. (1949) Experimental studies in rote-learning theory: VIII. Distributed practice of paired-associates with varying rates of presentation. *J. exp. Psychol.*, **39**, 714-718.
- HOVLAND, C. I. (1951) Human learning and retention. In Stevens, S. S. (Ed.) *Handbook of experimental psychology*. New York: Wiley.
- HOVLAND, C. I. (1952) A "communication analysis" of concept learning. *Psychol. Rev.*, **59**, 461-472.
- HOVLAND, C. I., AND KURTZ, K. H. (1951) Experimental studies in rote-learning theory: IX. Influence of work-decrement factors on verbal learning. *J. exp. Psychol.*, **42**, 265-272.
- HOVLAND, C. I., AND WEISS, W. (1953) Transmission of information concerning concepts through positive and negative instances. *J. exp. Psychol.*, **45**, 175-182.
- HULL, C. L. (1921) Quantitative aspects of the evolution of concepts: an experimental study. *Psychol. Monogr.*, **28**, No. 123.
- HULL, C. L. (1930) Knowledge and purpose as habit mechanisms. *Psychol. Rev.*, **37**, 511-525.
- HULL, C. L. (1931) Goal attraction and directing ideas conceived as habit phenomena. *Psychol. Rev.*, **38**, 487-506.
- HULL, C. L. (1932) The goal gradient hypothesis and maze learning. *Psychol. Rev.*, **39**, 25-43.
- HULL, C. L. (1933) The meaningfulness of 320 selected nonsense syllables. *Amer. J. Psychol.*, **45**, 730-734.
- HULL, C. L. (1935) The conflicting psychologies of learning—a way out. *Psychol. Rev.*, **42**, 491-516.
- HULL, C. L. (1943) *The principles of behavior*. New York: Appleton-Century-Crofts.
- HULL, C. L. (1947) Reactively heterogeneous compound trial and error learning with distributed trials and terminal reinforcement. *J. exp. Psychol.*, **37**, 118-135.
- HULL, C. L. (1948) Reactively heterogeneous compound trial and error learning with distributed trials and serial reinforcement. *J. exp. Psychol.*, **38**, 17-28.
- HULL, C. L. (1950) A primary social science law. *Sci. Monthly*, **71**, 221-228.
- HULL, C. L. (1951) *Essentials of behavior*. New Haven, Conn.: Yale Univer. Press.
- HULL, C. L. (1952) *A behavior system*. New Haven, Conn.: Yale Univer. Press.
- HULL, C. L., FELSINGER, J. M., GLADSTONE, A. I., AND YAMAGUCHI, H. G. (1947) A proposed quantification of habit strength. *Psychol. Rev.*, **54**, 237-254.
- HULL, C. L., HOVLAND, C. I., ROSS, R. T., HALL, M., PERKINS, D. T., AND FITCH, F. B. (1940) *Mathematico-deductive theory of rote learning*. New Haven, Conn.: Yale Univer. Press.

- HUMPHREY, G. (1933) *The nature of learning in its relation to the living system*. New York: Harcourt, Brace.
- HUMPHREYS, L. G. (1939) The effect of random alternation of reinforcement on the acquisition and extinction of conditioned eyelid reactions. *J. exp. Psychol.*, **25**, 141-158.
- HUNT, H. F., AND BRADY, J. V. (1955) Some effects of punishment and intercurrent "anxiety" on a simple operant. *J. comp. physiol. Psychol.*, **48**, 305-310.
- HUNT, J. McV. (1943) Motor behavior of rats in a conflict situation. Cited by Miller, N. E. (1944).
- HUNTER, I. M. L. (1952) An experimental investigation of the absolute and relative theories of transposition behaviour in children. *Brit. J. Psychol.*, **43**, 113-128.
- HUNTER, W. S. (1912) The delayed reaction in animals and children. *Behav. Monogr.*, **2**, No. 1.
- HUNTER, W. S. (1924) The symbolic process. *Psychol. Rev.*, **31**, 478-497.
- HUNTER, W. S. (1928) The behavior of raccoons in a double alternation temporal maze. *J. genet. Psychol.*, **35**, 374-388.
- HUNTER, W. S., AND BARTLETT, S. C. (1948) Double alternation behavior in young children. *J. exp. Psychol.*, **38**, 558-567.
- HURWITZ, H. M. B. (1954) Response-duration of lever pressing in the rat. *Quart. J. exp. Psychol.*, **6**, 62-71.
- HURWITZ, H. M. B. (1955) Response elimination without performance. *Quart. J. exp. Psychol.*, **7**, 1-7.
- INHELDER, B., AND PIAGET, J. (1955) *De la logique de l'enfant à la logique de l'adolescent*. Paris: Presses Univer. de France.
- IRION, A. L., AND GUSTAFSON, L. M. (1952) "Reminiscence" in bilateral transfer. *J. exp. Psychol.*, **43**, 321-323.
- JAHNKE, J. C., AND DUNCAN, C. P. (1956) Reminiscence and forgetting in motor learning after extended rest intervals. *J. exp. Psychol.*, **52**, 273-282.
- JAYNES, J. (1956) Imprinting: The interaction of learned and innate behavior: I. Development and generalization. *J. comp. physiol. Psychol.*, **49**, 201-206.
- JENKINS, J. G. (1933) Instruction as a factor in "incidental" learning. *Amer. J. Psychol.*, **45**, 471-477.
- JENKINS, W. O., MCFANN, H., AND CLAYTON, F. L. (1950) A methodological study of extinction following aperiodic and continuous reinforcement. *J. comp. physiol. Psychol.*, **43**, 155-167.
- JENKINS, W. O., AND POSTMAN, L. (1948) Isolation and "spread of effect" in serial learning. *Amer. J. Psychol.*, **61**, 214-221.
- JENKINS, W. O., AND SHEFFIELD, F. D. (1946) Rehearsal and guessing habits as sources of the "spread of effect." *J. exp. Psychol.*, **36**, 316-330.
- JENKINS, W. O., AND STANLEY, J. C. (1950) Partial reinforcement: a review and critique. *Psychol. Bull.*, **47**, 193-234.
- JENSEN, M. B., AND LEMOIRE, A. (1937) Ten experiments on whole and part learning. *J. educ. Psychol.*, **28**, 37-54.
- JOHDAL, K. (1955) Extinction as due to the changed direction of a psychological force. *J. exp. Psychol.*, **49**, 193-199.
- JOHNS, W. B. (1939) The growth of vocabulary among university students with some consideration of methods for fostering it. *J. exp. Educ.*, **8**, 89-102.
- JONES, E. I., AND BILODEAU, E. A. (1952) Differential transfer of training between motor tasks of different difficulty. *USAF Hum. Resour. Res. Cent. Bull.*, No. 52-35.
- JONES, H. E. (1930) The retention of conditioned emotional reactions in infancy. *J. genet. Psychol.*, **37**, 485-498.

- JONES, H. E. (1931) The conditioning of overt emotional responses. *J. educ. Psychol.*, **22**, 127-130.
- JUDSON, A. I., AND COFER, C. N. (1956) Reasoning as an associative process: I. "Direction" in a simple verbal problem. *Psychol. Rep.*, **2**, 469-476.
- JUDSON, A. I., COFER, C. N., AND GELFAND, S. (1956) Reasoning as an associative process: II. "Direction" in problem solving as a function of prior reinforcement of relevant responses. *Psychol. Rep.*, **2**, 501-507.
- KARN, H. W. (1938) The behavior of cats in a double alternation problem in the temporal maze. *J. comp. Psychol.*, **26**, 201-208.
- KATONA, G. *Organizing and memorizing*. New York: Columbia Univer. Press.
- KELLOGG, W. N., AND WALKER, E. L. (1938) "Ambiguous conditioning" a phenomenon of bilateral transfer. *J. comp. Psychol.*, **26**, 63-77.
- KELLOGG, W. N., AND WOLF, I. S. (1940) "Hypotheses" and "random activity" during the conditioning of dogs. *J. exp. Psychol.*, **26**, 588-601.
- KENDLER, H. H. (1945) Drive interaction: II. Experimental analysis of the role of drive in learning theory. *J. exp. Psychol.*, **35**, 188-198.
- KENDLER, T. S. (1950) An experimental investigation of transposition as a function of the difference between training and test stimuli. *J. exp. Psychol.*, **40**, 552-562.
- KIENTZLE, M. J. (1946) Properties of learning curves under varied distribution of practice. *J. exp. Psychol.*, **36**, 187-211.
- KIMBLE, G. A. (1947) Conditioning as a function of the time between conditioned and unconditioned stimuli. *J. exp. Psychol.*, **37**, 1-15.
- KIMBLE, G. A. (1949a) An experimental test of a two factor theory of inhibition. *J. exp. Psychol.*, **39**, 15-23.
- KIMBLE, G. A. (1949b) Performance and reminiscence in motor learning as a function of the degree of distribution of practice. *J. exp. Psychol.*, **39**, 500-510.
- KIMBLE, G. A. (1950) Evidence for the role of motivation in determining the amount of reminiscence in pursuit rotor learning. *J. exp. Psychol.*, **40**, 248-253.
- KIMBLE, G. A., AND BILODEAU, E. A. (1949) Work and rest as variables in cyclical motor learning. *J. exp. Psychol.*, **39**, 150-157.
- KIMBLE, G. A., AND KENDALL, J. W., JR. (1953) A comparison of two methods of producing experimental extinction. *J. exp. Psychol.*, **45**, 87-90.
- KIMBLE, G. A., MANN, L. I., AND DUFORT, R. H. (1955) Classical and instrumental eyelid conditioning. *J. exp. Psychol.*, **49**, 407-417.
- KITSON, H. D. (1922) A study of the output of workers under a particular wage incentive. *Univer. J. Bus.*, **1**, 54-68.
- KLEBANOFF, S. G. (1939) An experimental analysis of approach-approach and avoidance-avoidance conflict. New Haven, Conn., Dissertation, Yale Univer. Cited in Miller, N. E. (1944).
- KNIGHT, F. B. (1924) Transfer within a narrow mental function. *Elem. Sch. J.*, **24**, 780-788.
- KNÖPFELMACHER, F. (1953) Fixations, position stereotypes and their relation to the degree and pattern of stress. *Quart. J. exp. Psychol.*, **5**, 108-127, 150-157.
- KOCH, S. (1954) Clark L. Hull. In Estes, W. K., et al., *Modern learning theory*. New York: Appleton-Century-Crofts.
- KÖHLER, W. (1925) *The mentality of apes*. New York: Harcourt, Brace.
- KONORSKI, J. (1950) Mechanisms of learning. In *Symposia of the society for experimental biology*. No. 4.
- KRECHEVSKY, I. "Hypotheses" vs. "chance" in the pre-solution period in sensory discrimination learning. *Univer. Calif. Publ. Psychol.*, **6**, 27-44.

- KREUGER, W. C. F. (1929) The effect of overlearning on retention. *J. exp. Psychol.*, **12**, 71-78.
- KUENNE, M. R. (1946) Experimental investigation of the relation of language to transposition behavior in young children. *J. exp. Psychol.*, **36**, 471-490.
- KURTZ, A. K., WALTER, J. S., AND BRENNER, H. (1950) The effects of inserted questions and statements on film learning. *Penn. State Coll. Tech. Rep.*, SDC-297-16-1950.
- LAMBERT, W. W., LAMBERT, E. C., AND WATSON, P. D. (1953) Acquisition and extinction of an instrumental response sequence in the token-reward situation. *J. exp. Psychol.*, **45**, 321-326.
- LAMBERT, W. W., AND SOLOMON, R. L. (1952) Extinction of a running response as a function of distance of block point from the goal. *J. comp. physiol. Psychol.*, **45**, 269-279.
- LASHLEY, K. S. (1938) The mechanism of vision: XV. Preliminary studies of the rat's capacity for detail vision. *J. gen. Psychol.*, **18**, 123-193.
- LASHLEY, K. S. (1942) An examination of the "continuity theory" as applied to discriminative learning. *J. gen. Psychol.*, **26**, 241-265.
- LASHLEY, K. S., AND WADE, M. (1950) The Pavlovian theory of generalization. *Psychol. Rev.*, **53**, 72-87.
- LAUER, D. W., AND ESTES, W. K. (1955) Successive acquisition and extinction of a jumping habit in relation to schedule of reinforcement. *J. comp. physiol. Psychol.*, **48**, 8-13.
- LAWRENCE, D. H. (1949) Acquired distinctiveness of cues: I. Transfer between discriminations on the basis of familiarity with the stimulus. *J. exp. Psychol.*, **39**, 770-784.
- LAWRENCE, D. H. (1950) Acquired distinctiveness of cues: II. Selective association in a constant stimulus situation. *J. exp. Psychol.*, **40**, 175-188.
- LAWRENCE, D. H. (1952) The transfer of a discrimination along a continuum. *J. comp. physiol. Psychol.*, **45**, 511-516.
- LAWRENCE, D. H. (1954) The evaluation of training and transfer programs in terms of efficiency measures. *J. Psychol.*, **38**, 367-382.
- LAWRENCE, D. H. (1955) The applicability of generalization gradients to the transfer of a discrimination. *J. gen. Psychol.*, **52**, 37-48.
- LAWRENCE, D. H., AND DE RIVERA, J. (1954) Evidence for relational discrimination. *J. comp. physiol. Psychol.*, **47**, 465-471.
- LAWRENCE, D. H., AND MASON, W. A. (1955) Systematic behavior during discrimination reversal and change of dimension. *J. comp. physiol. Psychol.*, **48**, 267-271.
- LAZARUS, R. S., DEESE, J., AND HAMILTON, R. (1954) Anxiety and stress in learning: the role of intraserial duplication. *J. exp. Psychol.*, **47**, 111-114.
- LAZARUS, R. S., DEESE, J., AND OSLER, S. F. (1952) The effects of psychological stress upon performance. *Psychol. Bull.*, **49**, 293-317.
- LAZARUS, R. S., AND ERIKSEN, C. W. (1952) Psychological stress and its personality correlates: I. The effect of failure stress upon skilled performance. *J. exp. Psychol.*, **43**, 100-105.
- LAZARUS, R. S., YOUSEM, H., AND ARENBERG, D. (1953) Hunger and perception. *J. Pers.*, **21**, 312-329.
- LEVINE, J. M., AND MURPHY, G. (1943) The learning and forgetting of controversial material. *J. abnorm. soc. Psychol.*, **38**, 507-517.
- LEWIN, K. (1931) Environmental forces in child behavior and development. In Murchison, C. (Ed.) *A handbook of child psychology*. Worcester, Mass.: Clark Univer. Press.

- LEWIN, K. (1935) *A dynamic theory of personality*. New York: McGraw-Hill.
- LEWIS, D., SMITH, P. N., AND MCALLISTER, D. E. (1952) Retroactive facilitation and interference in performance on the modified two-hand coordinator. *J. exp. Psychol.*, **44**, 44-50.
- LIDDELL, H. S. (1944) Conditioned reflex method and experimental neurosis. In Hunt, J. McV. (Ed.) *Personality and the behavior disorders*, Vol. 1. New York: Ronald.
- LINDSLEY, D. B. (1946) A study of performance under speed stress. *OSRD Publ. Bull. No. 1838*. U. S. Dept. Commerce.
- LINDSLEY, D. B. (1951) Emotion. In Stevens, S. S. (Ed.) *Handbook of experimental psychology*. New York: Wiley.
- LITTMAN, R. A. (1949) Conditioned generalization of the galvanic skin reaction to tones. *J. exp. Psychol.*, **39**, 868-882.
- LLOYD, D. P. C. (1950) Post-tetanic potentiation of response in monosynaptic reflex pathways of the spinal cord. *J. gen. Physiol.*, **33**, 147-170.
- LOESS, H. B., AND DUNCAN, C. P. (1952) Human discrimination learning with simultaneous and successive presentation of stimuli. *J. exp. Psychol.*, **44**, 215-221.
- LOGAN, F. A. (1952) The role of delay of reinforcement in determining reaction potential. *J. exp. Psychol.*, **43**, 393-399.
- LORENZ, K. (1935) Der Kumpan in der Umwelt des Vogels. *J. Ornith.*, **83**, 137-213; 289-413.
- LORGE, I. (1930) Influence of regularly interpolated time intervals upon subsequent learning. *Teach. Coll. Contr. Educ.*, No. 438.
- LORGE, I., AND THORNDIKE, E. L. (1935) The influence of delay in the aftereffect of a connection. *J. exp. Psychol.*, **18**, 186-194.
- LOWELL, E. L. (1952) The effect of need for achievement on learning and speed of performance. *J. Psychol.*, **33**, 31-40.
- LUCHINS, A. S. (1942) Mechanization in problem solving. The effect of einstellung. *Psychol. Monogr.*, **54**, No. 248.
- LYON, D. O. (1917) *Memory and the learning process*. Baltimore, Md.: Warwick and York.
- MAATSCH, J. L., ADELMAN, H. M., AND DENNY, M. R. (1954) Effort and resistance to extinction of the bar-pressing response. *J. comp. physiol. Psychol.*, **47**, 47-50.
- MACCORQUODALE, K. (1947) An analysis of certain cues in delayed response. *J. comp. physiol. Psychol.*, **40**, 239-254.
- MACCORQUODALE, K., AND MEEHL, P. E. (1948) On a distinction between hypothetical constructs and intervening variables. *Psychol. Rev.*, **55**, 95-107.
- MACCORQUODALE, K., AND MEEHL, P. E. (1951) On the elimination of cul entries without obvious reinforcement. *J. comp. physiol. Psychol.*, **44**, 367-371.
- MACE, C. A. (1935) Incentives: some experimental studies. *Industr. Hlth. Res. Bd. Rep.*, No. 72.
- MAHER, W. B., AND WICKENS, D. D. (1954) Effect of differential quantity of reward on acquisition and performance of a maze habit. *J. comp. physiol. Psychol.*, **47**, 44-46.
- MAIER, N. R. F. (1929) Reasoning in white rats. *Comp. Psychol. Monogr.*, **6**, No. 29.
- MAIER, N. R. F. (1930) Reasoning in humans: I. On direction. *J. comp. Psychol.*, **10**, 115-143.
- MAIER, N. R. F. (1931) Reasoning in humans: II. The solution of a problem and its appearance in consciousness. *J. comp. Psychol.*, **12**, 181-194.
- MAIER, N. R. F. (1949) *Frustration*. New York: McGraw-Hill
- MALMO, R. B. (1954) Eccles' neurophysiological model of the conditioned reflex. *Canad. J. Psychol.*, **8**, 125-129.

- MALTZMAN, I., AND MORRISSETT, L. (1952) Different strengths of set in the solution of anagrams. *J. exp. Psychol.*, **44**, 242-246.
- MALTZMAN, I., AND MORRISSETT, L. (1953) Effects of task instructions on solution of different classes of anagrams. *J. exp. Psychol.*, **45**, 351-354.
- MANDLER, G. (1955) Associative frequency and associative prepotency as response measures of nonsense syllables. *Amer. J. Psychol.*, **68**, 662-665.
- MANDLER, G., AND SARASON, S. B. (1952) A study of anxiety and learning. *J. abnorm. soc. Psychol.*, **47**, 166-173.
- MARKS, M. R., AND JACK, O. (1952) Verbal context and memory span for meaningful material. *Amer. J. Psychol.*, **65**, 298-300.
- MARX, M. H., AND BUNCH, M. E. (1951) New gradients of error reinforcement in multiple-choice human learning. *J. exp. Psychol.*, **41**, 93-104.
- MASSERMAN, J. H. (1943) *Behavior and neurosis*. Chicago: Univer. of Chicago Press.
- MATHEWS, R. (1954) Recall as a function of number of classificatory categories. *J. exp. Psychol.*, **47**, 241-247.
- MCCLELLAND, D. C., ATKINSON, J. W., CLARK, R. W., AND LOWELL, E. L. (1953) *The achievement motive*. New York: Appleton-Century-Crofts.
- MCCORD, F. (1939) The delayed reaction and memory in rats: I. Length of delay. *J. comp. Psychol.*, **27**, 1-37.
- MCCRARY, J. W., JR., AND HUNTER, W. S. (1953) Serial position curves in verbal learning. *Science*, **117**, 131-134.
- MCCULLOCH, T. L., AND PRATT, T. G. (1934) A study of the presolution period in weight discrimination by white rats. *J. comp. Psychol.*, **18**, 271-290.
- MCGEOCH, J. A. (1930) The influence of associative value upon the difficulty of nonsense-syllable lists. *J. genet. Psychol.*, **37**, 421-426.
- MCGEOCH, J. A. (1936) The direction and extent of intra-serial association at recall. *Amer. J. Psychol.*, **48**, 221-245.
- MCGEOCH, J. A. (1942) *The psychology of human learning*. New York: Longmans.
- MCGEOCH, J. A., AND IRION, A. L. (1952) *The psychology of human learning*. (2d ed.) New York: Longmans.
- MCKINNEY, F., STROTHER, G. B., HINES, R. R., AND ALLEE, R. R. (1951) Experimental frustration in a group test situation. *J. abnorm. soc. Psychol.*, **46**, 316-323.
- MEEHL, P. E. (1950) On the circularity of the law of effect. *Psychol. Bull.*, **47**, 52-75.
- MELTON, A. W., AND IRWIN, J. MCQ. (1940) The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *Amer. J. Psychol.*, **53**, 173-203.
- MELTON, A. W., AND VON LACKUM, W. J. (1941) Retroactive and proactive inhibition in retention: evidence for a two factor theory of retroactive inhibition. *Amer. J. Psychol.*, **54**, 157-173.
- MEYER, D. R., BAHRICK, H. P., AND FITTS, P. M. (1953) Incentive, anxiety and the human blink rate. *J. exp. Psychol.*, **45**, 183-187.
- MILLER, G. A. (1951) *Language and communication*. New York: McGraw-Hill.
- MILLER, G. A. (1956a) Information theory and memory. *Sci. Amer.*, **195**, No. 2.
- MILLER, G. A. (1956b) The magical number seven; plus or minus two: some limits on our capacity for processing information. *Psychol. Rev.*, **63**, 81-97.
- MILLER, G. A., AND FRICK, F. C. (1949) Statistical behavioristics and sequences of responses. *Psychol. Rev.*, **56**, 311-324.
- MILLER, G. A., AND MCGILL, W. J. (1952) A statistical description of verbal learning. *Psychometrika*, **17**, 369-396.
- MILLER, G. A., AND SELFRIDGE, J. A. (1950) Verbal context and the recall of meaningful material. *Amer. J. Psychol.*, **63**, 176-185.

- MILLER, N. E. (1944) Experimental studies of conflict. In Hunt, J. McV. (Ed.) *Personality and the behavior disorders*, Vol. 1. New York: Ronald.
- MILLER, N. E. (1948) Studies of fear as an acquirable drive: I. Fear as motivation and fear-reduction as reinforcement in the learning of new responses. *J. exp. Psychol.*, **38**, 89-101.
- MILLER, N. E. (1951) Learnable drives and rewards. In Stevens, S. S. (Ed.) *Handbook of experimental psychology*. New York: Wiley.
- MILLER, N. E., BROWN, J. S., AND LIPOFSKY, H. (1943) Unpublished data. Cited in Miller, N. E. (1944).
- MILLER, N. E., AND DOLLARD, J. (1941) *Social learning and imitation*. New Haven, Conn.: Yale Univer. Press.
- MILLER, N. E., AND KESSEN, M. L. (1952) Reward effects of food via stomach fistula compared with those of food via mouth. *J. comp. physiol. Psychol.*, **45**, 555-564.
- MILLER, N. E., AND KRAELING, D. (1952) Displacement: greater generalization of approach than avoidance in generalized approach-avoidance conflict. *J. exp. Psychol.*, **43**, 217-221.
- MOELLER, G. (1954) The CS-UCS interval in GSR conditioning. *J. exp. Psychol.*, **48**, 162-166.
- MOLDAWSKY, S., AND MOLDAWSKY, P. C. (1952) Digit span as an anxiety indicator. *J. consult. Psychol.*, **16**, 115-118.
- MOLTZ, H. (1955) Latent extinction and the reduction of secondary reward value. *J. exp. Psychol.*, **49**, 395-400.
- MONTAGUE, E. K. (1953) The role of anxiety in serial rote learning. *J. exp. Psychol.*, **45**, 91-96.
- MONTGOMERY, K. C. (1951) The relation between exploratory behavior and spontaneous alternation in the white rat. *J. comp. physiol. Psychol.*, **44**, 582-589.
- MONTGOMERY, K. C. (1952) A test of two explanations of spontaneous alternation. *J. comp. physiol. Psychol.*, **45**, 287-293.
- MONTGOMERY, K. C. (1953) The effect of activity deprivation upon exploratory behavior. *J. comp. physiol. Psychol.*, **46**, 438-441.
- MONTGOMERY, K. C. (1954) The role of exploratory drive in learning. *J. comp. physiol. Psychol.*, **47**, 60-64.
- MORGAN, C. L. (1894) *An introduction to comparative psychology*. New York: Scott.
- MORGAN, R. L., AND UNDERWOOD, B. J. (1950) Proactive inhibition as a function of response similarity. *J. exp. Psychol.*, **40**, 592-603.
- MOWRER, O. H. (1947) On the dual nature of learning—a reinterpretation of "conditioning" and "problem-solving." *Harvard educ. Rev.*, **17**, 102-148.
- MOWRER, O. H. (1950) *Learning theory and personality dynamics*. New York: Ronald.
- MOWRER, O. H. (1953) Motivation and neuroses. In Jones, M. R. (Ed.) *Current theory and research in motivation, a symposium*. Lincoln, Nebr.: Univer. of Nebraska Press.
- MOWRER, O. H. (1954) Ego psychology, cybernetics and learning theory. In *Learning theory, personality theory, and clinical research* (the Kentucky symposium). New York: Wiley.
- MOWRER, O. H., AND AIKEN, E. G. (1954) Contiguity vs. drive-reduction in conditioned fear: temporal variations in conditioned and unconditioned stimulus. *Amer. J. Psychol.*, **67**, 26-38.
- MOWRER, O. H., AND JONES, H. M. (1945) Habit strength as a function of the pattern of reinforcement. *J. exp. Psychol.*, **35**, 293-311.

- MOWRER, O. H., AND LAMOREAUX, R. R. (1942) Avoidance conditioning and signal duration: a study of secondary motivation and reward. *Psychol. Monogr.*, **54**, No. 247.
- MOWRER, O. H., AND SOLOMON, L. N. (1954) Contiguity vs. drive-reduction in conditioned fear: the proximity and abruptness of drive-reduction. *Amer. J. Psychol.*, **67**, 15-25.
- MUELLER, C. G. (1950) Theoretical relationships among some measures of conditioning. *Proc. Nat. Acad. Sci. Wash.*, **36**, 123-130.
- MUNN, N. L. (1950) *Handbook of psychological research on the rat*. Boston, Mass.: Houghton Mifflin.
- MURRAY, E. J., AND MILLER, N. E. (1952) Displacement: steeper gradient of generalization of avoidance than approach with age of habit controlled. *J. exp. Psychol.*, **43**, 222-226.
- MYERS, A. K., AND MILLER, N. E. (1954) Failure to find a learned drive based on hunger: evidence for learning motivated by "exploration." *J. comp. physiol. Psychol.*, **47**, 428-440.
- NEWMAN, E. B. (1939) The effect of crowding of material on curves of forgetting. *Amer. J. Psychol.*, **52**, 601-609.
- NEWMAN, E. B., AND GERSTMAN, L. J. (1952) A new method for analyzing printed English. *J. exp. Psychol.*, **44**, 114-125.
- NIEMARK, E., AND SALTZMAN, I. J. (1953) Intentional and incidental learning with different rates of stimulus presentation. *Amer. J. Psychol.*, **66**, 618-621.
- NISSEN, N. W. (1950) Description of the learned response in discrimination behavior. *Psychol. Rev.*, **57**, 121-131.
- NISSEN, N. W., CHOW, K. L., AND SEMMES, J. (1951) Effects of restricted opportunity for tactual, kinesthetic and manipulative experience on the behavior of a chimpanzee. *Amer. J. Psychol.*, **14**, 485-507.
- NOBLE, C. E. (1952a) An analysis of meaning. *Psychol. Rev.*, **59**, 421-430.
- NOBLE, C. E. (1952b) The role of stimulus meaning (m) in serial verbal learning. *J. exp. Psychol.*, **43**, 437-446.
- NOBLE, C. E. (1955) Compound trial and error learning as a function of response availability (Nr). *J. exp. Psychol.*, **4**, 93-96.
- NOBLE, C. E. (1956) *Human trial and error learning*. Baton Rouge, La.: Human Learning Res. Laboratory, Louisiana State Univer.
- NOBLE, C. E., AND MCNEELY, D. A. (1957) The role of meaningfulness (m) in paired-associate verbal learning. *J. exp. Psychol.*, **53**, 16-22.
- NORTH, A. J. (1950) Improvement in successive discrimination reversals. *J. comp. physiol. Psychol.*, **43**, 442-460.
- O'KELLY, L. I., AND HEYER, A. W., JR. (1948) Studies in motivation and retention: I. Retention of a simple habit. *J. comp. physiol. Psychol.*, **41**, 466-478.
- OLDS, J. (1955) Physiological mechanisms of reward. In Jones, M. R. (Ed.) *Nebraska symposium on motivation*. Lincoln, Nebr.: Univer. of Nebraska Press.
- OLDS, J., AND MILNER, P. (1954) Positive reinforcement produced by electrical stimulation of the septal area and other regions of the rat brain. *J. comp. physiol. Psychol.*, **47**, 419-427.
- OSEAS, L., AND UNDERWOOD, B. J. (1952) Studies of distributed practice: V. Learning and retention of concepts. *J. exp. Psychol.*, **43**, 143-148.
- OSGOOD, C. E. (1946) Meaningful similarity and interference in learning. *J. exp. Psychol.*, **36**, 277-301.
- OSGOOD, C. E. (1949) The similarity paradox in human learning: a resolution. *Psychol. Rev.*, **56**, 132-143.

- OSGOOD, C. E. (1953) *Method and theory in experimental psychology*. New York: Oxford.
- OSLER, S. F. (1954) Intellectual performance as a function of two types of psychological stress. *J. exp. Psychol.*, **47**, 115-121.
- OVERMAN, J. R. (1930) An experimental study of the effect of the method of instruction on transfer in arithmetic. *Ele. Sch. J.*, **31**, 183-190.
- PAVLOV, I. P. (1927) *Conditioned reflexes*. (Trans. by G. V. Anrep.) New York: Oxford.
- PAVLOV, I. P. (1928) *Lectures on conditioned reflexes*. Vol. 1. (Trans. by W. H. Gantt.) New York: International Publishers.
- PERIN, C. T. (1942) Behavior potentiality as a joint function of the amount of training and the degree of hunger at the time of extinction. *J. exp. Psychol.*, **32**, 95-109.
- PERIN, C. T. (1943) The effect of delayed reinforcement upon the differentiation of bar responses in white rats. *J. exp. Psychol.*, **32**, 95-109.
- PERKINS, C. C. (1947) The relation of secondary reward to gradients of reinforcement. *J. exp. Psychol.*, **37**, 377-392.
- PIÉRON, H. (1920) Recherches comparatives sur la mémoire des formes et celle des chiffres. *Année psychol.*, **21**, 119-148.
- PORTER, J. M., JR. (1939) Experimental extinction as a function of the interval between successive non-reinforced elicitations. *J. gen. Psychol.*, **20**, 109-134.
- PORTER, L. W., AND DUNCAN, C. P. (1953) Negative transfer in verbal learning. *J. exp. Psychol.*, **46**, 61-64.
- POSTMAN, L. (1947) The history and present status of the law of effect. *Psychol. Bull.*, **46**, 489-563.
- POSTMAN, L., AND ADAMS, P. A. (1956) Studies in incidental learning: IV. The interaction of orienting task and stimulus materials. *J. exp. Psychol.*, **51**, 329-333.
- POSTMAN, L., ADAMS, P. A., AND BOHM, A. M. (1956) Studies in incidental learning: V. Recall for order and associative clustering. *J. exp. Psychol.*, **51**, 334-342.
- POSTMAN, L., ADAMS, P. A., AND PHILLIPS, L. W. (1955) Studies in incidental learning: II. The effects of association value and of the method of testing. *J. exp. Psychol.*, **49**, 1-10.
- POSTMAN, L., AND ALPER, T. G. (1946) Retroactive inhibition as a function of time of interpolation of the inhibitor between learning and recall. *Amer. J. Psychol.*, **59**, 439-449.
- POSTMAN, L., AND KAPLAN, H. L. (1947) Reaction time as a measure of retroactive inhibition. *J. exp. Psychol.*, **37**, 136-145.
- POSTMAN, L., AND TUMA, A. H. (1954) Latent learning in human subjects. *Amer. J. Psychol.*, **67**, 119-123.
- PRACHTL, H. F. R. (1953) Zur Physiologie der angeboren auslösenden Mechanism. *Behav.*, **5**, 32-50.
- RAMSEY, A. O., AND HESS, E. H. (1954) A laboratory approach to the study of imprinting. *Wilson Bull.*, **66**, 196-206.
- RAPAPORT, D. (1950) *Emotions and memory*. New York: International Publishers.
- REID, L. S. (1953) The development of noncontinuity behavior through continuity learning. *J. exp. Psychol.*, **46**, 107-112.
- RESTLE, F. (1955) A theory of discrimination learning. *Psychol. Rev.*, **62**, 11-19.
- REYNOLDS, B. (1945a) Extinction of a trace conditioned response as a function of the spacing of trials during the acquisition and extinction series. *J. exp. Psychol.*, **35**, 81-91.
- REYNOLDS, B. (1945b) A repetition of the Blodgett experiment on "latent" learning. *J. exp. Psychol.*, **35**, 504-516.

- REYNOLDS, B. (1948) The acquisition of a black-white discrimination habit under two levels of reinforcement. *J. exp. Psychol.*, **39**, 760-769.
- REYNOLDS, B. (1949) The relationship between strength of habit and the degree of drive present during acquisition. *J. exp. Psychol.*, **39**, 296-305.
- REYNOLDS, B. (1950a) Resistance to extinction as a function of the amount of reinforcement present during acquisition. *J. exp. Psychol.*, **40**, 46-52.
- REYNOLDS, B. (1950b) Acquisition of a simple spatial discrimination as a function of the amount of reinforcement. *J. exp. Psychol.*, **40**, 152-160.
- RICHTER, C. P. (1927) Animal behavior and internal drives. *Quart. Rev. Biol.*, **2**, 307-343.
- RIESEN, A. (1950) Arrested vision. *Sci. Amer.*, **183**, No. 2.
- RILEY, D. A., AND SHAPIRO, A. M. (1952) Alternation behavior as a function of effortfulness of task and distribution of trials. *J. comp. physiol. Psychol.*, **45**, 468-475.
- RIOPELLE, A. (1953) Transfer suppression and learning sets. *J. comp. physiol. Psychol.*, **46**, 108-114.
- RITCHIE, B. F., EBELING, E., AND ROTH, W. (1950) Evidence for continuity in the discrimination of vertical and horizontal patterns. *J. comp. physiol. Psychol.*, **43**, 168-180.
- ROBINSON, E. A. (1932) *Association theory today*. New York: Appleton-Century-Crofts.
- ROBINSON, F. P. (1946) *Effective study* (Rev. ed.). New York: Harper.
- ROCKWAY, M. R. (1953) Bilateral reminiscence in pursuit-rotor learning as a function of amount of first-hand practice and length of rest. *J. exp. Psychol.*, **46**, 337-344.
- ROHRER, J. H. (1947) Experimental extinction as a function of distribution of extinction trials and response strength. *J. exp. Psychol.*, **37**, 473-493.
- ROTHKOPF, E. Z., AND ZEAMAN, D. (1952) Some stimulus controls of alternation behavior. *J. Psychol.*, **34**, 235-255.
- ROTTER, J. B. (1954) *Social learning and clinical psychology*. Englewood Cliffs, N. J.: Prentice-Hall.
- RUBENSTEIN, H., AND ABORN, M. (1954) Immediate recall as a function of degree of organization and length of study period. *J. exp. Psychol.*, **48**, 146-152.
- RUGER, H. A. (1910) The psychology of efficiency: an experimental study of the process involved in the solution of mechanical puzzles and in the acquisition of skill in their manipulation. *Arch. Psychol.*, N.Y., **2**, No. 15.
- RUSSELL, W. A., AND STORMS, L. H. (1955) Implicit verbal chaining in paired-associate learning. *J. exp. Psychol.*, **49**, 287-293.
- SALTZMAN, I. J. (1949) Maze learning in the absence of primary reinforcement: a study of secondary reinforcement. *J. comp. physiol. Psychol.*, **42**, 161-173.
- SALTZMAN, I. J. (1953) The orienting task in incidental and intentional learning. *Amer. J. Psychol.*, **66**, 593-597.
- SALTZMAN, I. J. (1956) Comparisons of incidental and intentional learning with different orienting tasks. *Amer. J. Psychol.*, **69**, 274-277.
- SARASON, S. B., MANDLER, G., AND CRAIGHILL, P. G. (1952) The effect of differential instructions on anxiety and learning. *J. abnorm. soc. Psychol.*, **47**, 561-565.
- SCHLOSBERG, H. (1937) The relationship between success and the laws of conditioning. *Psychol. Rev.*, **44**, 379-394.
- SCHLOSBERG, H., AND KATZ, A. (1943) Double alternation lever-pressing in the white rat. *Amer. J. Psychol.*, **56**, 274-282.
- SCHRODER, H. M., AND ROTTER, J. B. (1952) Rigidity as learned behavior. *J. exp. Psychol.*, **44**, 141-150.

- SCHULZ, R. W. (1955) Generalization of serial position in rote serial learning. *J. exp. Psychol.*, **49**, 267-272.
- SEWARD, J. P. (1942) An experimental study of Guthrie's theory of reinforcement. *J. exp. Psychol.*, **30**, 247-256.
- SEWARD, J. P. (1948) The sign of a symbol: a reply to Professor Allport. *Psychol. Rev.*, **55**, 277-296.
- SEWARD, J. P. (1949) An experimental analysis of latent learning. *J. exp. Psychol.*, **39**, 177-186.
- SEWARD, J. P. (1951) Experimental evidence for the motivating function of reward. *Psychol. Bull.*, **48**, 130-149.
- SEWARD, J. P. (1953) How are motives learned? *Psychol. Rev.*, **60**, 99-110.
- SEWARD, J. P. (1954) Learning theory and identification: II. Role of punishment. *J. genet. Psychol.*, **84**, 201-210.
- SEWARD, J. P., AND LEVY, N. (1949) Sign learning as a factor in extinction. *J. exp. Psychol.*, **39**, 660-668.
- SEWARD, J. P., AND WELDON, R. J. (1953) Response latency as a function of change in delay of reward. *J. comp. physiol. Psychol.*, **46**, 184-189.
- SHARLOCK, D. P. (1954) The effects of a pre-extinction procedure on the extinction of a place and response performance in a T-maze. *J. exp. Psychol.*, **48**, 31-36.
- SHEFFIELD, F. D. (1949) "Spread of effect" without reward or learning. *J. exp. Psychol.*, **39**, 575-579.
- SHEFFIELD, F. D., AND CAMPBELL, B. A. (1954) The role of experience in the "spontaneous" activity of hungry rats. *J. comp. physiol. Psychol.*, **47**, 97-100.
- SHEFFIELD, F. D., AND JENKINS, W. O. (1952) Level of repetition in the "spread of effect." *J. exp. Psychol.*, **44**, 101-107.
- SHEFFIELD, F. D., AND ROBY, T. B. (1950) Reward value of nutritive and non-nutritive sweet taste. *J. comp. physiol. Psychol.*, **43**, 471-481.
- SHEFFIELD, V. F. (1949) Extinction as a function of partial reinforcement and distribution of practice. *J. exp. Psychol.*, **39**, 511-526.
- SHEFFIELD, V. F. (1950) Resistance to extinction as a function of the distribution of extinction trials. *J. exp. Psychol.*, **40**, 305-313.
- SHOBEN, E. J., JR. (1949) Psychotherapy as a problem in learning theory. *Psychol. Bull.*, **46**, 366-392.
- SIEGAL, A. I. (1953) Deprivation of visual form definition in the ring dove. *J. comp. physiol. Psychol.*, **46**, 115-119.
- SIEGAL, P. S., AND STEINBERG, M. (1949) Activity level as a function of hunger. *J. comp. physiol. Psychol.*, **42**, 413-416.
- SKINNER, B. F. (1938) *The Behavior of organisms*. New York: Appleton-Century-Crofts.
- SKINNER, B. F. (1950) Are theories of learning necessary? *Psychol. Bull.*, **57**, 193-216.
- SKINNER, B. F. (1954) The science of learning and the art of teaching. *Harvard educ. Rev.*, **24**, 86-97.
- SMALL, W. S. (1899) An experimental study of the mental processes of the rat. *Amer. J. Psychol.*, **11**, 133-164.
- SMALL, W. S. (1900) Experimental study of the mental processes of the rat. *Amer. J. Psychol.*, **12**, 206-239.
- SMITH, F. L. (1949) Generalization of antedating and perseverating tendencies in the linear maze. Baltimore, Md., Dissertation, The Johns Hopkins Univer.
- SMITH, M. H. (1949) Spread of effect is the spurious result of non-random response tendencies. *J. exp. Psychol.*, **39**, 355-368.

- SMITH, M. P. (1951) The stimulus trace gradient in visual discrimination learning. *J. comp. physiol. Psychol.*, **44**, 154-161.
- SMITH, O. A., AND GEIS, A. D. (1956) Comparative psychology in wildlife conservation. *Amer. Psychologist*, **11**, 183-187.
- SMOKE, K. L. (1932) An objective study of concept formation. *Psychol. Monogr.*, **42**, No. 191.
- SOLOMON, R. L. (1948) Effort and extinction rate: a confirmation. *J. comp. physiol. Psychol.*, **41**, 93-101.
- SOLOMON, R. L., KAMIN, L. J., AND WYNNE, L. C. (1953) Traumatic avoidance learning: the outcome of several extinction procedures with dogs. *J. abnorm. soc. Psychol.*, **48**, 291-302.
- SOLOMON, R. L., AND WYNNE, L. C. (1954) Traumatic avoidance learning: the principles of anxiety conservation and partial irreversibility. *Psychol. Rev.*, **61**, 353-385.
- SPENCE, K. W. (1932) The order of eliminating blinds in maze learning by the rat. *J. comp. Psychol.*, **14**, 9-27.
- SPENCE, K. W. (1936) The nature of discrimination learning in animals. *Psychol. Rev.*, **43**, 427-449.
- SPENCE, K. W. (1937a) The differential response in animals to stimuli varying within a single dimension. *Psychol. Rev.*, **44**, 430-444.
- SPENCE, K. W. (1937b) Analysis of formation of visual discrimination habits in the chimpanzee. *J. comp. Psychol.*, **23**, 77-100.
- SPENCE, K. W. (1940) Continuous versus non-continuous interpretations of discrimination learning. *Psychol. Rev.*, **47**, 271-288.
- SPENCE, K. W. (1942) The differential response in animals to stimuli varying within a single dimension. *J. exp. Psychol.*, **31**, 257-271.
- SPENCE, K. W. (1947) The role of secondary reinforcement in delayed reward learning. *Psychol. Rev.*, **54**, 1-8.
- SPENCE, K. W. (1952) The nature of the response in discrimination learning. *Psychol. Rev.*, **59**, 89-93.
- SPENCE, K. W. (1956) *Behavior theory and conditioning*. New Haven, Conn.: Yale Univer. Press.
- SPENCE, K. W., AND LIPPITT, R. (1946) An experimental test of the sign-Gestalt theory of trial and error learning. *J. exp. Psychol.*, **36**, 491-502.
- SPENCE, K. W., AND NORRIS, E. B. (1950) Eyelid conditioning as a function of the inter-trial interval. *J. exp. Psychol.*, **40**, 716-720.
- SPENCE, K. W., AND TAYLOR, J. A. (1951) Anxiety and strength of the UCS as determiners of the amount of eyelid conditioning. *J. exp. Psychol.*, **42**, 183-188.
- SPOONER, A., AND KELLOGG, W. N. (1947) The backward conditioning curve. *Amer. J. Psychol.*, **60**, 321-334.
- SPROW, A. J. (1947) Reactively homogeneous compound trial and error learning with distributed trials and terminal reinforcement. *J. exp. Psychol.*, **37**, 197-213.
- STANLEY, W. C. (1952) Extinction as a function of the spacing of extinction trials. *J. exp. Psychol.*, **43**, 249-259.
- STANLEY, W. C., AND ROWE, M. I. (1954) Extinction by omission of food as a function of goal box confinement. *J. exp. Psychol.*, **48**, 271-274.
- STELLAR, E. (1954) The physiology of motivation. *Psychol. Rev.*, **61**, 5-22.
- STEPHENS, J. M. (1934) Further notes on punishment and reward. *J. genet. Psychol.*, **44**, 464-472.
- STEPHENS, J. M. (1942) Expectancy vs. effect-substitution as a general principle of reinforcement. *Psychol. Rev.*, **49**, 102-116.

- STEPHENS, J. M. (1951) Spontaneous schooling: a neglected feature in theories of education. *Sch. & Soc.*, **73**, 332-341.
- STEPHENS, J. M. (1956) Non-deliberative factors underlying the phenomenon of schooling. *Educ. Theory*, **6**, 26-34.
- STEVENSON, H. W., ISCOE, I., AND MCCONNELL, C. (1955) A developmental study of transposition. *J. exp. Psychol.*, **49**, 278-280.
- STEWART, C. N., AND WARREN, J. M. (1957) The behavior of cats on the double-alternation problem. *J. comp. physiol. Psychol.*, **50**, 26-28.
- STONE, G. R. (1948) A note on Postman's review of the literature on the law of effect. *Psychol. Bull.*, **45**, 151-160.
- STRANGE, J. R. (1950) Latent learning under conditions of high motivation. *J. comp. physiol. Psychol.*, **43**, 194-197.
- STRANGE, J. R. (1954) The effect of an irrelevant drive on the reaction tendency specific to another drive. *J. gen. Psychol.*, **51**, 31-40.
- STRASSBURGER, R. C. (1950) Resistance to extinction of a conditioned operant as related to drive level at reinforcement. *J. exp. Psychol.*, **40**, 473-487.
- STROUD, J. B. (1940) Experiments on learning in school situations. *Psychol. Bull.*, **37**, 777-807.
- SWIFT, E. J. (1918) *Psychology and the day's work: a study in the application of psychology to daily life*. New York: Scribner's.
- TAFT, R. (1954) Selective recall and memory distortion of favorable and unfavorable material. *J. abnorm. soc. Psychol.*, **49**, 23-28.
- TAYLOR, D. W. (1943) Learning telegraphic code. *Psychol. Bull.*, **40**, 461-487.
- TAYLOR, J. A. (1951) The relationship of anxiety to the conditioned eyelid response. *J. exp. Psychol.*, **41**, 81-92.
- TAYLOR, J. A. (1956) Drive theory and manifest anxiety. *Psychol. Bull.*, **53**, 303-320.
- TAYLOR, J. A., AND SPENCE, K. W. (1952) The relationship of anxiety level to performance in serial learning. *J. exp. Psychol.*, **44**, 61-64.
- TEICHNER, W. H. (1952) Experimental extinction as a function of the intertrial intervals during conditioning and extinction. *J. exp. Psychol.*, **44**, 170-178.
- THISTLETHWAITE, D. (1950) Attitude and structure as factors in the distortion of reasoning. *J. abnorm. soc. Psychol.*, **45**, 442-458.
- THISTLETHWAITE, D. (1951) A critical review of latent learning and related experiments. *Psychol. Bull.*, **48**, 97-130.
- THOMPSON, M. E., AND THOMPSON, J. E. (1949) Reactive inhibition as a factor in maze learning: II. The role of reactive inhibition in studies of place learning versus response learning. *J. exp. Psychol.*, **39**, 883-891.
- THOMPSON, W. R., AND SOLOMON, L. N. (1954) Spontaneous pattern discrimination in the rat. *J. comp. physiol. Psychol.*, **47**, 104-107.
- THORNDIKE, E. L. (1898) Animal intelligence. *Psychol. Monogr.*, **2**, No. 8.
- THORNDIKE, E. L. (1915) Watson's "Behavior." *J. anim. Behav.*, **5**, 452-467.
- THORNDIKE, E. L. (1923) The influence of first-year Latin upon the ability to read English. *Sch. & Soc.*, **17**, 165-168.
- THORNDIKE, E. L. (1932a) Reward and punishment in animal learning. *Comp. Psychol. Monogr.*, **8**, No. 39.
- THORNDIKE, E. L. (1932b) *The fundamentals of learning*. New York: Teach. Coll.
- THORNDIKE, E. L. (1933) An experimental study of rewards. *Teach. Coll. Contr. Educ.*, No. 580.
- THORNDIKE, E. L. (1935) *The psychology of wants, interests and attitudes*. New York: Appleton-Century-Crofts.

- THORNDIKE, E. L., AND LORGE, I. (1944) *The teacher's word book of 30,000 words*. New York: Teach. Coll.
- THORNDIKE, E. L., AND RUGER, G. J. (1923) The effect of first-year Latin upon knowledge of English words of Latin derivation. *Sch. & Soc.*, **18**, 260-270.
- THORNDIKE, E. L., AND WOODWORTH, R. S. (1901) The influence of improvement in one mental function upon the efficiency of other functions. *Psychol. Rev.*, **8**, 247-267, 384-395, 553-564.
- THORPE, W. H. (1956) *Learning and instinct in animals*. Cambridge, Mass.: Harvard Univer. Press.
- THUNE, L. E., AND UNDERWOOD, B. J. (1943) Retroactive inhibition as a function of degree of interpolated learning. *J. exp. Psychol.*, **32**, 185-200.
- THURSTONE, L. L. (1930a) The learning function. *J. gen. Psychol.*, **3**, 469-493.
- THURSTONE, L. L. (1930b) The relation between learning time and length of task. *Psychol. Rev.*, **37**, 44-53.
- TINBERGEN, N. (1951) *The study of instinct*. New York: Oxford.
- TINBERGEN, N. (1953) *Social behaviour in animals*. London: Methuen.
- TOLMAN, E. C. (1932) *Purposive behavior in animals and men*. New York: Appleton-Century-Crofts.
- TOLMAN, E. C. (1938) The law of effect. *Psychol. Rev.*, **45**, 165-203.
- TOLMAN, E. C. (1948) Cognitive maps in rats and men. *Psychol. Rev.*, **55**, 189-208.
- TOLMAN, E. C. (1949) There is more than one kind of learning. *Psychol. Rev.*, **56**, 144-155.
- TOLMAN, E. C. (1951) *Collected papers in psychology*. Berkeley, Calif.: Univer. of Calif. Press.
- TOLMAN, E. C., AND HONZIK, C. H. (1930) Introduction and removal of reward and maze performance in rats. *Univ. Calif. Publ. Psychol.*, **4**, 257-275.
- TOLMAN, E. C., RITCHIE, B. F., AND KALISH, D. (1946) Studies in spatial learning vs. response learning. *J. exp. Psychol.*, **36**, 221-229.
- TUFTS COLLEGE INSTITUTE OF APPLIED EXPERIMENTAL PSYCHOLOGY (1952) Handbook of human engineering data. Special Devices Center, *Tech. Rep. No. SDC 199-1-2*.
- ULMER, G. (1939) Teaching geometry to cultivate reflective thinking: an experimental study with 1239 high school pupils. *J. exp. Educ.*, **8**, 18-25.
- UNDERWOOD, B. J. (1945) The effect of successive interpolations on retroactive and proactive inhibition. *Psychol. Monogr.*, **59**, No. 273.
- UNDERWOOD, B. J. (1948) Retroactive and proactive inhibition after 5 and 48 hours. *J. exp. Psychol.*, **38**, 29-38.
- UNDERWOOD, B. J. (1951) Studies of distributed practice: II. Learning and retention of paired-adjective lists with two levels of intra-list similarity. *J. exp. Psychol.*, **42**, 153-161.
- UNDERWOOD, B. J. (1953a) Studies of distributed practice: X. The influence of intralist similarity on learning and retention of serial adjective lists. *J. exp. Psychol.*, **45**, 253-259.
- UNDERWOOD, B. J. (1953b) Studies of distributed practice: XI. An attempt to resolve conflicting facts on retention of serial lists. *J. exp. Psychol.*, **45**, 355-359.
- UNDERWOOD, B. J., (1954a) Intralist similarity in verbal learning and retention. *Psychol. Rev.*, **61**, 160-166.
- UNDERWOOD, B. J. (1954b) Speed of learning and amount retained: a consideration of methodology. *Psychol. Bull.*, **51**, 276-282.
- UNDERWOOD, B. J. (1957) Interference and forgetting. *Psychol. Rev.*, **64**, 49-60.

- UNDERWOOD, B. J., AND GOAD, D. (1951) Studies of distributed practice: I. The influence of intra-list similarity in serial learning. *J. exp. Psychol.*, **42**, 125-134.
- UNDERWOOD, B. J., AND RICHARDSON, J. (1956) The influence of meaningfulness, intra-list similarity and serial position on retention. *J. exp. Psychol.*, **52**, 119-126.
- VERPLANCK, W. S. (1956) The operant conditioning of human motor behavior. *Psychol. Bull.*, **53**, 70-83.
- VERPLANCK, W. S., AND HAYES, J. R. (1953) Eating and drinking as a function of maintenance schedule. *J. comp. physiol. Psychol.*, **46**, 327-333.
- WALK, R. D. (1952) Effect of discrimination reversal on human discrimination learning. *J. exp. Psychol.*, **44**, 410-419.
- WARD, L. B. (1937) Reminiscence and rote learning. *Psychol. Monogr.*, **49**, No. 220.
- WARREN, J. M. (1954) An analysis of the formation of visual discriminative habits of Rhesus monkeys. *Amer. J. Psychol.*, **67**, 517-520.
- WASSERMAN, H. N. (1951) The effect of motivation and amount of pre-rest practice upon inhibitory potential in motor learning. *J. exp. Psychol.*, **42**, 162-172.
- WATSON, J. B. (1907) Kinaesthetic and organic sensations: their role in the reactions of the white rat. *Psychol. Rev. Monogr.*, **8**, No. 33.
- WATSON, J. B. (1914) *Behavior: an introduction to comparative psychology*. New York: Holt.
- WATSON, J. B. (1924) *Psychology from the standpoint of a behaviorist*. (2d ed.) Philadelphia, Pa.: Lippincott.
- WATSON, J. B., AND RAYNOR, R. (1920) Conditioned emotional reactions. *J. exp. Psychol.*, **3**, 1-14.
- WEBB, W. B. (1949) The motivational aspect of an irrelevant drive in the behavior of the white rat. *J. exp. Psychol.*, **39**, 1-14.
- WEBB, W. B. (1950) A test of "relational" vs. "specific stimulus" learning in discrimination problems. *J. comp. physiol. Psychol.*, **43**, 70-72.
- WEINSTOCK, S. (1954) Resistance to extinction of a running response following partial reinforcement under widely spaced trials. *J. comp. physiol. Psychol.*, **47**, 318-322.
- WEISS, P., AND BITTERMAN, M. E. (1951) Response selection in discriminative learning. *Psychol. Rev.*, **58**, 185-195.
- WELCH, G. B., AND BURNETT, C. T. (1924) Is primacy a factor in association formation? *Amer. J. Psychol.*, **35**, 396-401.
- WERTHEIMER, M. (1945) *Productive thinking*. New York: Harper.
- WHITFIELD, J. W. (1951) An experiment on problem solving. *Quart. J. exp. Psychol.*, **3**, 184-196.
- WHITING, J. W. M., AND CHILD, I. L. (1953) *Child training and personality: a cross-cultural study*. New Haven, Conn.: Yale Univer. Press.
- WHITING, J. W. M., AND MOWRER, O. H. (1943) Habit progression and regression—a laboratory study of some factors relevant to human socialization. *J. comp. Psychol.*, **36**, 229-253.
- WIKE, E. L. (1953) Extinction of a partially and continuously reinforced response with and without a rewarded alternative. *J. exp. Psychol.*, **46**, 255-260.
- WIKE, E. L., AND MCNAMARA, H. J. (1955) A quest for the generalized conditioned reinforcer. *Psychol. Rep.*, **1**, 83-91.
- WILCOX, M. J. (1917) Does the study of high school Latin improve high school English? *Sch. & Soc.*, **6**, 56-60.
- WILLIAMS, S. B. (1938) Resistance to extinction as a function of the number of reinforcements. *J. exp. Psychol.*, **23**, 506-522.
- WILLIAMS, S. B. (1943) A note on approach-approach conflict in rats. *J. comp. Psychol.*, **35**, 269-274.

- WOLFE, J. B. (1934) The effect of delayed reward upon learning in the white rat. *J. comp. Psychol.*, **17**, 1-21.
- WOLFE, J. B. (1936) Effectiveness of token rewards for chimpanzees. *Comp. Psychol. Monogr.*, **12**, No. 60.
- WOLFLE, H. M. (1930) Time factors in conditioned finger withdrawal. *J. gen. Psychol.*, **7**, 372-378.
- WOLFLE, H. M. (1931) Conditioning as a function of the interval between the conditioned and original stimulus. *J. gen. Psychol.*, **7**, 80-103.
- WOODBURY, C. B. (1950) Double, triple and quadruple repetition in the white rat. *J. comp. physiol. Psychol.*, **43**, 490-502.
- WOODROW, H. (1927) The effect of type of training upon transference. *J. educ. Psychol.*, **18**, 159-172.
- WOODWORTH, R. S., AND SCHLOSBERG, H. (1954) *Experimental psychology*. (Rev. ed.) New York: Holt.
- YERKES, R. M. (1943) *Chimpanzees*. New Haven, Conn.: Yale Univer. Press.
- YOUNG, P. T. (1936) *Motivation of behavior*. New York: Wiley.
- YOUNG, P. T., AND SHUFORD, E. H., JR. (1954) Intensity, duration and repetition of hedonic processes as related to acquisition of motives. *J. comp. physiol. Psychol.*, **47**, 298-305.
- YOUNG, R. K., AND UNDERWOOD, B. J. (1954) Transfer in verbal materials with dissimilar stimuli and response similarity varied. *J. exp. Psychol.*, **47**, 153-159.
- YUM, K. S. (1931) An experimental test of the law of assimilation. *J. exp. Psychol.*, **14**, 68-82.
- ZEAMAN, D. (1949) Response latency as a function of the amount of reinforcement. *J. exp. Psychol.*, **39**, 466-483.
- ZEAMAN, D., AND HOUSE, B. J. (1951) The growth and decay of reactive inhibition as measured by alternation behavior. *J. exp. Psychol.*, **41**, 177-186.
- ZELLER, A. F. (1950) An experimental analogue of repression: II. The effect of individual failure and success on memory measured by relearning. *J. exp. Psychol.*, **40**, 411-422.
- ZELLER, A. F. (1951) An experimental analogue of repression: III. The effect of individual failure and success on memory measured by recall. *J. exp. Psychol.*, **42**, 32-38.
- ZIRKLE, G. A. (1946) Success and failure in serial learning: II. Isolation and the Thorndike effect. *J. exp. Psychol.*, **36**, 302-315.

NAME INDEX

- Aborn, M., 179-180, 270, 333, 348
 Adams, D. K., 278, 333
 Adams, P. A., 121, 348
 Adamson, R. E., 287-288, 333
 Adelman, H. M., 50, 61, 333, 344
 Aiken, E. G., 131, 346
 Alberts, E., 85, 333
 Allee, R. R., 149, 345
 Alper, T. G., 260, 348
 Ammons, R. B., 195, 333
 Anderson, E., 109, 333
 Arenberg, D., 115, 343
 Arnold, W. J., 153, 333
 Arnoult, M. D., 252, 333
 Atkinson, J. W., 119, 345
 Attneave, F., 250-251, 333
 Atwater, S. K., 259, 333
 Auble, D., 328-329, 333
 Auld, F., 324, 336
 Austin, G. A., 296-299, 335

 Bahrnick, H. P., 117, 345
 Baker, R. A., 89, 99, 333
 Baron, M. R., 97, 334
 Bartlett, F. C., 249-253, 333
 Bartlett, S. C., 282, 341
 Bass, M. J., 76, 333
 Battig, W. F., 229-230, 334
 Bauer, F. J., 91-92, 334
 Baum, M. H., 293-295, 334
 Bilodeau, E. A., 97, 100, 230, 334, 341, 342
 Birch, H. G., 277-278
 Bitterman, M. E., 89-90, 117, 334, 354
 Blodgett, H. C., 25-27, 154, 334
 Bohm, A. M., 121, 348

 Bolden, L., 53, 59, 335
 Book, W. F., 186, 334
 Bousfield, W. A., 176, 244-246, 334
 Brady, J. V., 140*n.*, 341
 Brandauer, C. M., 102, 334
 Brenner, H., 202, 343
 Briggs, G. E., 260, 265, 334
 Briggs, L. J., 243, 334
 Brogden, W. J., 127, 204-206, 334
 Brown, J. S., 97, 109, 132-133, 334, 346
 Bruce, R. W., 222, 257, 334
 Bruner, J. S., 296-299, 335
 Brunswick, E., 71-73, 335
 Bryan, W. L., 184-186, 335
 Buchwald, A. M., 118, 335
 Bugelski, B. R., 56, 257, 335
 Bunch, M. E., 195, 201, 335, 345
 Burke, C. J., 22, 38-39, 306, 337
 Burnett, C. T., 175, 354
 Burt, H. E., 271, 335
 Bush, R. R., 22, 72, 78, 306, 335
 Buss, A. H., 164, 335
 Buss, E., 164, 335
 Butler, R. A., 40, 53, 335
 Buxton, C. E., 27, 335

 Cadwallader, T. C., 257, 335
 Calvin, A. D., 53, 59, 335
 Campbell, B. A., 99, 108, 335, 350
 Carpenter, J. A., 107, 336
 Carper, J. W., 107, 335
 Carr, H. A., 155, 335
 Child, I. L., 325, 335, 354
 Chow, K. L., 95, 347
 Clark, R. W., 119, 345
 Clayton, F. L., 67, 341

- Clifford, B., 53, 59, 335
 Clifford, L. T., 53, 59, 335
 Coate, W. B., 89, 334
 Cofer, C. N., 212, 285-286, 290, 335, 336, 342
 Cohen, B. H., 176, 245, 334
 Cole, L. E., 218, 336
 Cotterman, T. E., 284, 336
 Cowles, J. T., 109, 336
 Coyer, R. A., 56, 335
 Craighill, P. G., 118, 349
 Crawford, M. P., 192, 336
 Crawley, M. E., 216, 338
 Crespi, L. P., 24, 336
 Culler, E. A., 127, 334
- Darwin, C., 33
 Dattman, P. E., 293, 336
 Deese, J., 56-57, 106-107, 117, 148-149, 174-176, 179, 207, 225-226, 250-252, 258, 268, 336, 343
 Denny, M. R., 50, 344
 DeRivera, J., 86, 343
 Dethier, V. G., 320, 336
 Dinsmoor, J. A., 102, 140, 336
 Dollard, J., 35, 145-147, 323-324, 336, 346
 Dufort, R. H., 62, 128*n.*, 336, 342
 Duncan, C. P., 96, 188, 193, 222, 259, 336, 341, 344, 348
 Duncker, K., 286, 289, 337
 Dykman, R. A., 131, 337
- Ebbinghaus, E., 165, 169-170, 211, 237, 241-243, 266, 310, 337
 Ebeling, E., 88, 349
 Eccles, J. C., 316-317, 337
 Ehrenfreund, D., 85, 88, 333, 337
 Ellis, D. S., 191, 337
 Ellson, D. G., 51, 337
 Eriksen, C. W., 117, 149, 337, 343
 Estes, W. K., 16, 22, 38-39, 59-63, 72, 139-140, 184, 306, 308, 337
- Farber, I. E., 117, 337
 Feldman, R. S., 146, 337
 Felsing, J. M., 304, 340
 Finan, J. L., 106, 337
- Finger, F. A., 99, 337
 Fitch, F. B., 171, 340
 Fitts, P. M., 117, 345
 Foley, J. P., Jr., 290, 336
 Foster, H., 216, 338
 Frandsen, A., 233, 337
 French, E. G., 120, 337
 French, R. S., 252, 337
 Freud, S., 269-270, 337
 Frick, F. C., 153, 161-163, 283, 337, 338, 345
 Furchtgott, E., 25, 338
- Gagne, R. M., 16, 19, 49, 216, 338
 Galanter, E., 96, 338
 Gantt, W. H., 131, 143, 337, 338
 Gates, A. I., 202, 338
 Gatling, F., 88-89, 338
 Geis, A. D., 97, 351
 Gerstman, L. J., 177, 347
 Gibson, E. J., 97, 221, 229, 257, 338
 Gladstone, A. E., 159, 304, 338, 340
 Glanzer, M., 58, 62, 338
 Glaze, J. A., 165, 338
 Gleitman, H., 55, 59, 338
 Goad, D., 191, 207, 354
 Godbeer, E., 134, 338
 Goodnow, J. J., 296-299, 335
 Goodwin, W. R., 92, 338
 Goss, A. E., 229, 338
 Graham, C. H., 16, 19, 338
 Grant, D. A., 68, 72, 76, 338
 Greenspoon, J., 164, 338
 Grice, G. R., 32, 78-79, 88-90, 338, 339
 Grindley, G. C., 23, 339
 Gulliksen, H., 22*n.*, 339
 Gustafson, L. M., 194, 341
 Guthrie, E. R., 37-39, 59-63, 322, 339
 Guttman, N., 23, 339
- Hake, H. W., 72, 338
 Hall, M., 171, 340
 Hamilton, G. V., 161
 Hamilton, R., 207, 343
 Hardman, G. W., Jr., 225-226, 258, 268, 336
 Harlow, H. F., 23, 93, 103, 228, 279, 284, 335, 339
 Harter, N., 184-186, 335

- Harvey, J., 53, 59, 335
Hayes, J. R., 102, 354
Heathers, G. L., 58, 339
Hebb, D. O., 80, 94-95, 261, 314-316, 339
Heidbreder, E., 292-295, 339
Hess, E. H., 319, 348
Heyer, E. A., Jr., 107, 347
Hilgard, E. R., 8*n.*, 165, 302, 339
Hillman, B., 107, 339
Hinde, R. A., 62, 339
Hines, R. R., 149, 345
Hobbs, N., 198, 339
Holland, G. A., 195, 339
Holtzman, W. H., 117, 334
Homme, L. E., 60*n.*, 339
Honzik, C. H., 62, 353
Hornseth, J. P., 72, 338
Horton, G. P., 38, 339
House, B. J., 58, 355
Hovland, C. I., 16, 20-21, 49, 53, 76, 80, 171-172, 191-192, 295-296, 340
Hull, C. L., 22, 28-29, 34-36, 40, 53-59, 76, 82-83, 102-104, 156-160, 165, 171-174, 195, 232, 292, 303-305, 333, 340
Humphrey, G., 8, 318, 340
Humphreys, L. G., 66, 341
Hunt, H. F., 140*n.*, 340
Hunt, J. McV., 134, 341
Hunter, I. M. L., 85, 341
Hunter, W. S., 107, 172-173, 280-282, 339, 341, 345
Hurwitz, H. M. B., 56, 150, 341

Inhelder, B., 284, 341
Irion, A. L., 194, 242, 341, 345
Irwin, J. McQ., 264, 345
Iscoe, I., 85, 352
Israel, H., 293, 336

Jack, O., 179, 345
Jahnke, J. C., 193, 341
Jaynes, J., 319, 341
Jenkins, J. G., 120, 341
Jenkins, W. O., 66-67, 201, 341, 350
Jensen, M. B., 203, 341
Johdai, K., 62, 341
Johns, W. B., 234, 341

Jones, E. I., 230, 341
Jones, H. E., 125-126, 341, 342
Jones, H. M., 49, 70, 346
Judson, A. I., 286, 342

Kalish, D., 95, 353
Kamin, L. J., 130, 350
Kaplan, H. L., 263, 348
Karn, H. W., 282, 342
Katona, G., 312-313, 342
Katz, A., 153, 283*n.*, 349
Kaufman, R. A., 175-176, 179, 336
Keenan, J., 149, 336
Kellogg, W. N., 43-44, 222, 342, 351
Kendall, J. W., Jr., 60, 342
Kendler, H. H., 106, 342
Kendler, T. S., 85, 342
Kessen, M. L., 35, 346
Kientzle, M. J., 189-190, 342
Kimble, G. A., 43, 60, 107, 128*n.*, 189-190, 193, 336, 339, 342
Kitson, H. D., 116, 342
Klebenoff, S. G., 134, 342
Knight, F. B., 226, 342
Knöpfelmacher, F., 146, 342
Koch, S., 55, 303, 342
Köhler, W., 83-84, 275-276, 313, 342
Konorski, J., 8*n.*, 342
Kraeling, D., 108, 134, 335, 346
Krechevsky, I., 87, 342
Kresse, F. H., 174, 336
Kreuger, W. C. F., 243, 343
Kuenne, M. R., 85, 343
Kurtz, A. K., 202, 343
Kurtz, K. H., 192, 340

Lambert, E. C., 61, 343
Lambert, W. W., 61, 343
Lamoreaux, R. R., 129, 347
Lashley, K. S., 79-81, 89, 343
Lauer, D. W., 60, 343
Lawrence, D. H., 86, 91-93, 217, 230, 235, 334, 338, 343
Lazarus, R. S., 115, 117, 148-149, 207, 336, 343
Lemoire, A., 203, 341
Levine, J. M., 253-254, 343
Lewin, K., 131-135, 343, 344
Lewis, D., 259, 344

- Liddell, H. S., 143, 344
 Lindsley, D. B., 123, 149, 344
 Lipman, E. A., 127, 334
 Lipofsky, H., 132, 346
 Lippitt, R. A., 106, 351
 Littman, R. A., 76, 344
 Lloyd, D. P. C., 316, 344
 Locke, J., 309-310
 Loess, H. B., 96, 344
 Logan, F. A., 29, 344
 Lorenz, K., 319, 344
 Lorge, I., 166, 188, 199, 344, 352
 Lowell, E. L., 119, 344, 345
 Luchins, A. S., 288, 344
 Lyon, D. O., 208-212, 344
- Maatsch, J. L., 50, 61, 333, 344
 McAllister, D. E., 259, 344
 McClearn, G. E., 103, 339
 McClelland, D. C., 119, 345
 McConnell, C., 85, 352
 McCord, F., 281, 345
 MacCorquodale, K., 26, 281, 344
 McCrary, J. W., Jr., 172-173, 345
 McCulloch, T. L., 88, 345
 McCutchan, K., 154, 334
 Mace, C. A., 116, 344
 McFann, H., 67, 341
 Macfarlane, F. A., 156, 344
 McGeoch, J. A., 170-171, 207, 242, 255, 310, 345
 McGill, W. J., 308, 345
 McNamera, H. J., 110*n.*, 354
 McNeely, D. A., 207, 347
 Magsdick, W. K., 195, 335
 Maher, W. B., 25, 344
 Maier, N. R. F., 145-146, 278, 288, 344
 Malmo, R. B., 317, 344
 Maltzman, I., 289, 345
 Mandler, G., 118, 165, 345, 349
 Mann, L. I., 128*n.*, 342
 Marder, V. J., 252, 336
 Marks, M. R., 179, 345
 Marquis, D. M., 8*n.*, 339
 Marx, M. H., 201, 345
 Mason, W. A., 92, 343
 Masserman, J. H., 141, 143-144, 345
 Mathews, R., 246, 345
 Mech, E. V., 328-329, 333
 Meehl, P. E., 18, 26, 344
- Melton, A. W., 165, 258, 264, 345
 Meumann, E., 211
 Meyer, D. R., 23, 117, 284, 336, 339, 345
 Miller, G. A., 161-163, 177-179, 206, 246-248, 283, 308, 338, 345
 Miller, N. E., 35, 110-112, 132-135, 145-147, 323-324, 336, 346, 347
 Milner, P., 36, 347
 Moeller, G., 43-45, 346
 Moldawsky, P. C., 149, 346
 Moldawsky, S., 149, 346
 Moltz, H., 56-57, 346
 Montague, E. K., 117, 346
 Montgomery, K. C., 58, 103, 346
 Morgan, C. L., 274, 346
 Morgan, R. L., 257, 346
 Morrissett, L., 289, 345
 Mosteller, F., 22, 72, 78, 306, 335
 Mowrer, O. H., 8*n.*, 13, 41, 49, 63, 70, 113, 129-131, 140, 324-325, 346, 347
 Mueller, C. G., 16, 347
 Munn, N. L., 160
 Murphy, G., 253-254, 343
 Murray, E. J., 135, 347
 Myers, A. K., 112, 347
- Nachmias, J., 55, 59, 338
 Neimark, E., 121, 347
 Neisser, U., 55, 59, 338
 Newman, E. B., 177, 243, 347
 Nissen, N. W., 90, 95, 347
 Noble, C. E., 165, 204-207, 252, 347
 Norris, E. B., 49, 351
 North, A. J., 89, 347
- O'Kelly, L. I., 107, 347
 Olds, J., 36, 347
 Oseas, L., 191, 347
 Osgood, C. E., 223-226, 257, 285*n.*, 292, 347, 348
 Osler, S. F., 148-149, 343, 348
 Overman, J. H., 234, 348
- Pavlov, I. P., 8-11, 30-31, 42-43, 49-52, 55, 75, 143, 348
 Perin, C. T., 27-29, 32, 48, 104, 348
 Perkins, C. C., 32, 348

- Perkins, D. T., 171, 340
 Phillips, L. W., 121, 341
 Piaget, J., 284, 341
 Piéron, H., 251, 348
 Porter, J. M., Jr., 49, 348
 Porter, L. W., 222, 348
 Postman, L., 121-122, 137, 201, 260, 263, 341, 348
 Prachtl, H. F. R., 58, 348
 Pratt, T. G., 88, 345
- Rall, W., 316, 337
 Ramsey, A. O., 319, 348
 Raynor, R., 125, 354
 Reed, H. B., 243, 334
 Reid, L. S., 92, 348
 Restle, F., 306, 348
 Reynolds, B., 24-26, 49, 106, 348, 349
 Richardson, J., 207, 354
 Richter, C. P., 99, 349
 Riesen, A., 94, 349
 Riley, D. A., 58, 349
 Riopelle, A., 93, 349
 Ritchie, B. F., 88, 95, 349, 353
 Robinson, E. A., 310, 349
 Robinson, F. P., 202, 349
 Roby, T. D., 36, 350
 Rockway, M. R., 194, 348
 Rogers, W. A., 57, 335
 Rohrer, J. H., 49, 349
 Ross, B. M., 68, 338
 Ross, R. T., 171, 340
 Roth, W., 88, 349
 Rothkopf, E. Z., 58, 349
 Rotter, J. B., 289, 326-327, 349
 Rowe, M. I., 61, 351
 Rubenstein, H., 179-180, 333, 349
 Rubin, R. D., 25, 338
 Ruger, H. A., 218, 279, 349, 352
 Russell, W. A., 171, 349
- Saltz, E., 78-79, 339
 Saltzman, I. J., 31, 121, 199, 347, 349
 Sarason, S. B., 118, 345, 349
 Schiller, J. J., 76, 338
 Schipper, L. M., 68, 338
 Schlosberg, H., 8*n.*, 90*n.*, 153, 217, 283*n.*, 349, 355
 Schmidt, R. E., 204-206, 334
- Schroder, H. M., 286, 349
 Schulz, R. W., 171, 350
 Selfridge, J. A., 177-179, 345
 Semmes, J., 95, 347
 Seward, J. P., 27, 29, 39, 56, 107, 109, 112, 281, 325-326, 350
 Shapiro, A. M., 58, 349
 Sharlock, D. P., 56, 350
 Shaw, W. A., 96, 338
 Sheffield, F. D., 36, 99, 201, 335, 341, 350
 Sheffield, V. F., 49, 350
 Shoben, E. J., Jr., 323, 350
 Shuford, E. H., Jr., 36, 355
 Siegal, A. I., 94, 350
 Siegal, P. S., 99, 350
 Silva, J. G., 176, 334
 Skinner, B. F., 8*n.*, 9-12, 16-17, 29, 47-48, 60-61, 64-68, 99-101, 138-139, 150, 322, 329, 350
 Small, W. S., 154, 350
 Smith, F. L., 159, 350
 Smith, M. H., 201, 350
 Smith, M. P., 32, 351
 Smith, O. A., 97, 351
 Smith, P. N., 259, 344
 Smith, S., 248
 Smoke, K. L., 292-293, 295, 350
 Solomon, L. N., 103, 131, 347, 352
 Solomon, R. L., 50, 61, 130-131, 141, 350
 Spence, K. W., 32, 49, 82-85, 89-90, 106, 117-118, 157, 305, 337, 350
 Spooner, A., 43-45, 351
 Sprow, A. J., 159, 351
 Stanley, J. C., 66, 341
 Stanley, W. C., 49, 61, 351
 Steinberg, M., 99, 350
 Stellar, E., 123, 351
 Stephens, J. M., 8*n.*, 137, 329-330, 351, 352
 Stevenson, H. W., 85, 352
 Stewart, C. N., 284, 352
 Stone, G. R., 137, 352
 Storms, L. H., 171, 349
 Strange, J. R., 100, 102, 352
 Strassburger, R. C., 106, 352
 Straughan, J. N., 72, 308, 337
 Strother, G. B., 149, 345
 Stroud, J. B., 218, 352
 Swift, E. J., 186, 352

- Taft, R., 254, 352
 Taylor, D. W., 186, 287-288, 333, 352
 Taylor, J. A., 116-118, 351, 352
 Teichner, W. H., 49, 352
 Thistlethwaite, D., 106, 290, 352
 Thompson, J. E., 96, 352
 Thompson, M. E., 96, 352
 Thompson, W. R., 103, 352
 Thorndike, E. L., 34, 136-138, 157, 166, 197, 199-201, 218, 274-275, 278, 327-328, 344, 352, 353
 Thorpe, W. H., 318-321, 353
 Thune, L. E., 259, 353
 Thurstone, L. L., 182-183, 209, 353
 Tinbergen, N., 75, 97, 353
 Tolman, E. C., 8, 25, 30, 40, 62, 95-96, 353
 Tuma, A. H., 122, 348

 Ulmer, G., 232, 353
 Underwood, B. J., 191, 193, 196, 207, 225, 229, 243-244, 257-259, 265-267, 310, 337, 346, 347, 353, 354, 355

 Verplanck, W. S., 102, 164, 354
 von Lackum, W. J., 258, 345

 Wade, M., 79, 87, 343
 Walk, R. D., 96, 354
 Walker, E. L., 222, 342
 Walter, J. S., 202, 343
 Ward, L. B., 173, 354
 Warren, J. M., 93, 284, 352, 354
 Wasserman, H. N., 196, 354

 Watson, J. B., 125, 154-157, 335, 354
 Watson, P. D., 61, 343
 Webb, W. B., 85, 102, 354
 Weiner, M., 164, 335
 Weinstock, S., 68, 354
 Weise, P., 90, 354
 Weiss, W., 296, 340
 Welch, G. B., 175, 354
 Weldon, R. J., 29, 350
 Wertheimer, M., 313, 354
 White, A. M., 324, 336
 Whitehorn, J. C., 131, 337
 Whitfield, J. W., 299*n.*, 354
 Whiting, J. W. M., 63, 140, 325, 354
 Wickens, D. D., 25, 284, 336, 344
 Wike, E. L., 67, 110*n.*, 354
 Williams, S. B., 48, 104, 134, 354
 Windecker, E., 97
 Wolf, I. S., 222, 342
 Wolfe, J. B., 29, 31, 109, 355
 Wolfe, H. M., 43-44, 355
 Woodbury, C. B., 159, 355
 Woodrow, H., 228, 355
 Woodworth, R. S., 90*n.*, 216, 218, 327, 352, 355
 Wynne, L. C., 130-131, 141, 350

 Yamaguchi, H. G., 118, 304, 335, 340
 Yerkes, R. M., 276-277, 355
 Young, P. T., 36, 115, 355
 Young, R. K., 225, 355
 Yousem, H., 115, 343
 Yum, K. S., 97, 220, 355

 Zeaman, D., 24, 58, 349, 355
 Zeller, A. F., 270, 355
 Zirkle, G. A., 200, 355

SUBJECT INDEX

- Ability and transfer, 234-235
- Acquired distinctiveness of cues, 91-93
- Activity, 99-100
- Alternation, double, 282-283
- Amount to be learned, 208-212
- Anticipatory goal response, fractional, 157-158
- Anxiety, 116-119, 135
 - fear as drive, 110-111
- Association theory, 36-40, 309-311
 - (*See also* Reinforcement, theories of)
- Association value and difficulty, 206-207
- Associations, backward, 169
 - remote, in verbal learning, 169-172
- Attainment of concepts (*see* Concept formation)
- Autocorrelation, 163
 - (*See also* Information theory)
- Avoidance conditioning, 127-131

- Backward associations, 169
- Backward conditioning, 43-44
- Behavior theory (*see* Learning theory)

- Categorization, 291-293
- Classical conditioning, 12-13, 41, 75-78
 - emotion in, 125-127
- Clustering, 244-246
- Competition theory, extinction, 59-61
 - forgetting, 261-269
- Concept formation, 291-300
 - positive and negative instances in, 295-296
 - strategies in, 297-300
- Conditioned inhibition, 54-55
 - (*See also* Extinction, theories of)
- Conditioning, avoidance, 127-131
 - (*See also* Punishment)
 - backward, 43-44
 - classical, 12-13, 41, 75-78, 125-127
 - emotion and, 125-127, 139-140
 - instrumental, 12-13, 41
 - temporal factors in, 42-45
 - two-factor theory of, 127-131
 - two types, 6-13, 41
- Conflict, 131-135
 - approach-approach, 134
 - approach-avoidance, 133
 - avoidance-avoidance, 133-134
- Contiguity (*see* Association theory)
- Continuity theory, 82-91

- Delayed reaction, 280-281
- Derived lists, method of, 169-170
- Differentiation and transfer, 229-230
 - (*See also* Discriminative learning)
- Difficulty, task, 204-208, 230-231
 - and association value, 206-207
- Discipline, formal, 217-218
- Discrimination, 74-97
- Discriminative learning, 81-94
 - acquired distinctiveness of cues, 91-93
 - learning sets in, 93-94
 - theories of, 82-91
 - algebraic summation, 82-86
 - continuity, 82-86
 - noncontinuity, 87-88
 - (*See also* Perceptual learning)
- Disinhibition, 50-51
- Distributed practice, theories of, 194-197

- Distribution of practice, 187-197
 extinction and, 48-49
 permanence of effects, 191-193
 rest periods in, 189-190
 in skills, 188-194
 task differences and, 190-191
 transfer and, 194
 Disuse, principle of, 255
 Double alternation, 282-283
 Drive (*see* Motivation)
- Education, and learning, 327-330
 and transfer of training, 217-219, 231-235
 Effect, principle of, 197-203
 spread of, 199-201
 (*See also* Reinforcement)
 Effort (*see* Distribution of practice; Extinction)
 Emotion, 123-149
 classical conditioning of, 125-127
 conflict and, 131-135
 punishment and, 135-141
 two-factor theory and, 127-131
 (*See also* Motivation)
 Ethology, 317-321
 Expectancy, 68-69
 Experimental designs in transfer, 214-216
 Experimental neurosis, 141-147
 in frustration, 145-146
 studies in, 142-143
 theories of, 144-146
 Extinction, 46-63
 distribution of trials, 48-49
 effort in, 49-50
 generalization of, 80
 partial reinforcement and, 66-68
 resistance to, 47-48, 66-68
 responding in, 55-57
 theories of, 52-63
 competition, 59-61
 frustration, 61
 incentive loss, 61-62
 response-produced inhibition, 53-59
- Fear as drive, 110-111
 Fixed-interval reinforcement, 64-65
 Fixed-ratio reinforcement, 65
 Forgetting, 255-272
 competition theory, 261-269
 proactive inhibition in, 256-269
 repression in, 269-271
 retroactive inhibition in, 256-269
 selective, 253-255
 stimulus-response relationships in, 257-258
 theories of, 260-269
 two-factor theory of, 265-269
 (*See also* Proactive inhibition; Retention; Retroactive inhibition)
 Formal discipline, 217-218
 Fractional anticipatory goal response, 157-158
 Free recall (*see* Recall)
 Frustration, 61, 145-146
 experimental neurosis in, 145-146
 Frustration theory in extinction, 61
 Functional fixedness, 286-288
- Generalization, 75-80, 158-160, 219-226
 extinction, 80
 in maze learning, 158-160
 response, 221-223
 stimulus, 75-80
 conditioning, classical, 75-80
 instrumental, 78-80
 in transfer, 219-221
 in transfer, 219-226
 Gestalt theory, 311-313
 Gradient, goal, 157-158
 of reinforcement, 32
 (*See also* Reinforcement)
- Habituation, 317-318
 Hierarchy of habits, 184-187
 Human behavior (*see* Motivation)
 Hypotheses in problem solving, 285-286, 289
 continuity theory, 82-91
 Hypothetical constructs, 6
- Identical elements, 218-219
 Immediate memory span, 246-248
 Imprinting, 317-319
 Incentives (*see* Motivation; Reinforcement)

- Incidental learning, 120-122
- Information in human learning, 197-203
- Information theory, animal behavior and, 161-163
 - recording and memory span, 246-248
 - statistical behavioristics, 161-163
 - in verbal learning, 176-180
- Insight, 275-279
- Instrumental conditioning, 12-13, 41
- Interference (*see* Competition theory)

- Knowledge of results, 197-203

- Latent learning, 25-27, 105-106
- Latent period, 19
- Learning, discriminative (*see* Discriminative learning)
 - education and, 327-330
 - equation for, 22
 - human, information in, 197-203
 - incidental, 120-122
 - latent, 25-27, 105-106
 - maze (*see* Maze learning)
 - and motivation, 105-108
 - paired-associate, 167-169
 - perceptual, 94-96
 - versus performance, 23-24, 29-30, 34-35
 - reinforcement theories and, 34-35
 - time of reinforcement and, 29-30
 - place, 95-96
 - in psychotherapy, 323-325
 - school, transfer in, 231-235
 - and socialization, 325-326
 - types of, 6-13
 - verbal (*see* Verbal learning)
- Learning how to learn, 227-228
- Learning curves, 21-23, 181-187
 - plateaus in, 184-187
 - theoretical, 21-23, 182-183
 - types of, 181-187
- Learning sets, 93-94, 227-228
- Learning theory, 302-313
 - association, 309-311
 - gestalt, 311-313
 - neurophysiological, 314-317
 - quantitative, 302-308
 - statistical, 305-308
 - (*See also* Reinforcement)
- Length of task, 208-212
- Linear maze, 158-160

- Magnitude of response, 19-20
- Maze learning, 151-160
 - linear maze, 158-160
 - theories of, 154-160
- Mediating responses, 286, 289
 - (*See also* Recoding)
- Memory (*see* Forgetting; Retention)
- Memory drum, 166
- Memory span, 246-248
- Memory trace, 260-261
- Motivation, 98-122
 - human behavior and, 113-122
 - experimental studies, 115-120
 - incentive value, 115-116
 - inventory studies, 116-119
 - origin, 114
 - projective tests in, 119-120
 - incidental learning and, 120-122
 - learning and, 105-108
 - performance and, 99-104
 - irrelevant needs, 102-103
 - secondary, 108-113
 - fear, 110-111
 - theories of, 112-113
 - token-reward studies of, 109-110
 - stress and, 147-148
- Need reduction, 35-36
- Needs (*see* Motivation)
- Negative transfer (*see* Proactive inhibition; Retroactive inhibition; Transfer of training)
- Neurophysiological theories of learning, 314-317
- Neurosis (*see* Experimental neurosis)
- Noncontinuity theory, 82-91
- Nonsense syllables, 165

- Operant behavior, 12
- Operational definitions, 5

- Paired-associate learning, 167-169
- Partial reinforcement, choice behavior and, 71-73

- Partial reinforcement,
 effect of, on extinction, 66–68
 on learning, 65–66
 fixed-interval, 64–65
 fixed ratio, 65
 theories of, 68–73
 discrimination, 69–70
 expectancy, 68–69
 response, 70–71
 variable-ratio, 65
 (*See also* Reinforcement)
- Perceptual learning, 94–96
 (*See also* Discriminative learning)
- Perceptual memory, 249–252
- Performance versus learning (*see* Learning)
- Periodic reinforcement (*see* Partial reinforcement)
- Personality theory, learning in, 321–327
- Place learning, 95–96
- Plateaus, 184–187
- Potentiation, 316–317
- Practice, conditions of, 187–203
 distribution of (*see* Distribution of practice)
- Proactive inhibition, 256–269
 competition theory of, 261–269
 and degree of learning, 259–260
 in forgetting, 256–269
 and retroactive inhibition, 258–259
 stimulus-response relationships in, 257–258
 theories of, 260–269
 and two-factor theory, 265–269
 (*See also* Transfer of training)
- Problem solving, 273–300
 in animals, 273–284
 concept formation in, 291–300
 fixation in, 286–290
 human, 284–300
 hypotheses in, 285–286, 289
 insightful, 275–279
 symbolic behavior in, 279–284
 trial and error, 274–275
- Psychotherapy, learning in, 323–325
- Punishment, 135–141
 in animal behavior, 138–141
 avoidance conditioning, 127–131
 conditioned emotional responses and, 139–140
 and emotion, 135–141
- Punishment,
 Thorndikian experiments in, 136–138
 traumatic, 141
- Quantitative theory, 303–309
- Reactive inhibition, 53–54
 (*See also* Extinction, theories of)
- Recall, free, 175–179
 measure of retention, 239
 number of categories in, 245–246
 serial position in, 175–176
 statistical organization in, 177–179
- Recoding, 246–248
- Reflex reserve, 47–48
- Reinforcement, 15–45
 amount of, 23–27
 definition of, 15–16
 delay of, 27–30
 example of, 11
 fixed-interval, 64–65
 fixed-ratio, 65
 gradient of, 32
 (*See also* Gradient)
 number of, 19–21
 partial (*see* Partial reinforcement)
 schedules of, 64–65
 secondary, 30–32
 and time of, 31–32
 theories of, 33–41
 association, 37–41
 contiguity, 37–41
 need-reduction, 35–36
 time of, 27–30
 trans-situational, 18
 variable-ratio, 65
 (*See also* Effect)
- Remote associations, 169–172
- Repression, 269–271
- Resistance to extinction, 47–48, 66–68
- Respondent behavior, 12
- Response, magnitude of, 19–20
 versus place learning, 95–96
 strength of, 16, 19–23
- Response generalization, 221–223
- Rest periods, 189–190
- Retention, 237–254
 clustering in, 244–246
 course of, 241–242

- Retention,
 dynamic factors in, 252-255
 measurement of, 237-241
 perceptual, 249-252
 recoding, 246-248
 reorganization in, 244-247
 speed of learning and, 243-244
 (*See also* Forgetting)
- Retroactive inhibition, 256-269
 competition theory of, 261-269
 and degree of learning, 259-260
 in forgetting, 256-269
 and proactive inhibition, 258-259
 stimulus-response relationships in,
 257-258
 theories of, 260-269
 and two-factor theory, 265-269
 (*See also* Transfer of training)
- Reward (*see* Effect; Reinforcement)
- Rote learning (*see* Verbal learning)
- Satiation, stimulus, 57-58
- Savings, 239-240
- Schematizing, 249-252
- School learning, transfer in, 231-235
 (*See also* Education)
- Secondary motivation (*see* Motivation)
- Secondary reinforcement (*see* Reinforcement)
- Serial learning, 150-180
 in animals, 151-163
 human (*see* Verbal learning)
 mazes and, 151-160
 theories of, 154-160
 verbal (*see* Verbal learning)
- Serial-position effect, 172-176
- Similarity, difficulty and, 207-208
 in transfer, 219-226
 (*See also* Generalization)
- Socialization and learning, 325-326
- Spontaneous recovery, 50-51
- Spread of effect, 199-201
- Statistical behavioristics, 161-163
- Statistical learning theory, 305-308
- Stimulus generalization (*see* Generalization)
- Stimulus satiation, 57-58
- Strength of response, 16, 19-23
- Stress, 147-149
- Symbolic behavior, 279-284
- Task difficulty (*see* Difficulty)
- Task length, 208-212
- Task variables, 203-212
- Temporal factors in conditioning, 42-45
- Thinking (*see* Problem solving)
- Token reward, 109-110
 (*See also* Reinforcement, secondary)
- Trace conditioned responses, 42-43
- Transfer of training, 213-235
 and ability, 234-235
 and differentiation, 229-230
 and difficulty, 230-231
 in distribution of practice, 194
 education and, 217-219, 231-235
 experimental designs in, 214-216
 generalization in, 219-226
 identical elements in, 218-219
 learning how to learn, 227-228
 measurement of, 213-216
 stimulus generalization in, 219-221
 stimulus-response relationships in,
 219-226
 theory of, 217-227
 training for, 232-234
 transfer surface, 223-226
 (*See also* Proactive inhibition; Retroactive inhibition)
- Transposition, 83-86
- Trans-situational reinforcement, 18
- Traumatic punishment, 141
- Trial and error, 161-162, 274-275
 vicarious, 91*n*.
- Two-factor theory, of conditioning, 127-131
 of forgetting, 265-269
- Uncertainty (*see* Information theory)
- Variable-ratio reinforcement, 65
- Verbal learning, 163-180
 remote associations in, 169-172
 rote methods of, 164-169
 paired associates, 167-169
 serial, 166-167
 serial-position effect in, 172-176
 statistical organization in, 176-180
 (*See also* Forgetting; Retention; Transfer of training)
- Vicarious trial and error, 91*n*.

30.3.5